

Biophysics: Searching for Principles

William Bialek*

(Dated: September 18, 2011)

This is a draft, not complete but hopefully not embarrassing either. I would be very happy to receive feedback at wbialek@princeton.edu. Please note the caveats in the section “About this draft,” which includes an explanation for the [red ellipses](#).

CONTENTS

Introduction	1
A. About our subject	1
Looking back	3
Looking forward	4
B. About this book	6
C. About this draft	8
Acknowledgments	8
I. Photon counting in vision	11
A. Posing the problem	11
B. Single molecule dynamics	29
C. Dynamics of biochemical networks	38
D. The first synapse, and beyond	49
E. Perspectives	59
II. Noise isn't negligible	63
A. Molecular fluctuations and chemical reactions	63
B. Molecule counting	79
C. More about noise in perception	101
D. Proofreading and active noise reduction	118
E. Perspectives	134
III. No fine tuning	135
A. Sequence ensembles	136
B. Ion channels and neuronal dynamics	153
C. The states of cells	166
D. Long time scales in neural networks	182
E. Perspectives	191
IV. Efficient representation	195
A. Entropy and information	195
B. Does biology care about bits?	218
C. Optimizing information flow	233
D. Gathering information and making models	247
E. Perspectives	258
V. Outlook	261
A. Appendix	263
1. Poisson processes	263
2. Correlations, power spectra and all that	270
3. Electronic transition in large molecules	276
4. Cooperativity	281
5. X-ray diffraction and biomolecular structure	286
6. Berg and Purcell, revisited	290
7. Dimensionality reduction	298

8. Maximum entropy	298
9. Measuring information transmission	304

INTRODUCTION

Like all authors, I hope that this book will find wide readership. At the same time, I believe that good books are intensely personal objects. As readers, we have our favorite books, and this is an emotional statement, laden with context.¹ Similarly, writers bring not just their knowledge and their technical skill to the creation of a book, but also their personalities. In writing something which might be used as a textbook, I feel a responsibility to provide a fair view of the field. But I won't apologize for giving you *my* view, which surely is not a consensus view. Indeed, perhaps by the time there is a clear consensus the field won't be quite as much fun.

A. About our subject

When a PhD student in Physics picks up a textbook about elementary particle physics, or cosmology, or condensed matter, there is little doubt about what will be found inside the covers. There are questions, perhaps, about the level and style of presentation, or about the emphasis given to different subfields, but the overall topic is clear. The situation is very different for books or courses that attempt to bring the intellectual style of physics to bear on the phenomena of life. The problem is not just in how we teach, but also in how we do research. The community of physicists interested in biological problems is incredibly diverse, it spills over into more amorphously defined interdisciplinary communities, and individual physicists often are more connected to biologists working on the same system than they are to physicists asking the same conceptual question in other

*William Bialek is the John Archibald Wheeler/Battelle Professor in Physics, and a member of the multidisciplinary Lewis-Sigler Institute for Integrative Genomics, at Princeton University. He also serves as Visiting Presidential Professor of Physics at the Graduate Center of the City University of New York.

¹ The book which gave me my first taste of real quantum mechanics has a special place in my library, even though it isn't a book I would recommend to my students. Translated from the Russian, it looks like it was typed rather than typeset. An important part of the story is that I found it for sale on a remainder table in a department store. It must have been the only quantum mechanics book ever sold by the *Emporium*.

systems. None of this is necessarily good or bad, but it can be terribly confusing for students.

Ours is not a new subject, but over its long history, “biophysics” or “biological physics” has come to mean many different things to different communities.² At the same time, for many physicists today, biophysics remains new, and perhaps a bit foreign. There is an excitement to working in a new field, and I hope to capture this excitement. Yet our excitement, and that of our students, sometimes is tempered by serious concerns, which can be summarized by naive questions: Where is the boundary between physics and biology? Is biophysics really physics, or just the application of methods from physics to the problems of biology? My biologist friends tell me that ‘theoretical biology’ is nonsense, so what would theoretical physicists be doing if they got interested in this field? In the interaction between physics and biology, what happens to chemistry? How much biology do I need to know in order to make progress? Why do physicists and biologists seem to be speaking such different languages? Can I be interested in biological problems and still be a physicist, or do I have to become a biologist? Although there has been much progress over the last decade, I still hear students (and colleagues) asking these questions, and so it seems worth a few pages to place the subject of this book into context.³ The discussion will start by reacting to the history of our subject, but by the end I hope to outline a view of the field which stands on its own as a guide to what we would like to accomplish, both on the time scale of working through this book and on the longer time scale of our research agendas [not quite sure about that last phrase, but want to say something in this spirit].

There is an old saying that “physics is what physicists do.” This doesn’t sound very helpful, but it may be getting at an important point. Academic disciplines have a choice to define themselves either by their objects of study or by their style of inquiry. Physics (at its best, I would like to think) is firmly in the second camp. Physicists make it their business to ask certain kinds of

questions about Nature, and to seek certain kinds of answers. “Thinking like a physicist” means something, and we are proud to do it; it is this, above all else, that we try to convey to our students. We are the intellectual heirs of Galileo, taking seriously his evocative claim that the book of Nature is written in the language of mathematics.

Biology surely is defined by the objects of study—if it’s not alive, biologists aren’t interested. The style of inquiry may change, from studies of animal behavior and anatomy to genetics and molecular structure, but the objects remain the same. It is especially important for physicists to appreciate the vastness of the enterprise that is labeled ‘biology,’ and the tremendous divisions within biology itself. A geneticist, for example, studying the dynamics of regulatory networks in a simple organism such as yeast, may know absolutely nothing about the dynamics of neural networks for the regulation of movement in higher organisms, and vice versa. Not only is biology defined by the objects of study, but the subfields of biology are similarly defined, so that networks of neurons and networks of genes are different subjects.

Differences in our view of the scientific enterprise translate rather directly into different educational structures. In physics, we (try to) teach principles and derive the predictions for particular examples. In biology, teaching proceeds (mostly) from example to example. Although physics has subfields, to a remarkable extent the physics community clings to the romantic notion that Physics is one subject. Not only is the book of Nature written in the language of mathematics, but there is only one book, and we expect that if we really grasped its content it could be summarized in very few pages. Where does biophysics fit into this view of the world?

There is something different about life, something that we recognize immediately as distinguishing the animate from the inanimate. But we no longer believe that there is a fundamental “life force” that animates a lump of inert stuff. Similarly, there is no motive force which causes superfluid helium to crawl up the sides of a container and escape, or which causes electrical current in a superconducting loop to flow forever; the phenomena of superfluidity and superconductivity emerge as startling consequences of well known interactions among electrons and nuclei, interactions which usually have much more mundane consequences. As physicists studying the phenomena of life, we thus are not searching for a new force of Nature. Rather we are trying to understand how the same forces that usually cause carbon based materials to look like rocks or sludge can, under some conditions, cause material to organize itself and walk (or swim or fly) out of the laboratory. What is special about the state of matter that we call life? How does it come to be this way? Different generations of physicists have approached these mysteries in different ways.

² The use of these two different words is also problematic. I think that, roughly speaking, “biophysics” can be used by people who think of themselves either as physicists or biologists, while “biological physics” is an attempt to carve out a subfield of physics, distinct from biology. The difficulty is that neither word really points to a set of questions that everyone can agree upon. So, we need to dig in.

³ The intellectual questions about biophysics and its relation to the larger, separate, activities of physics and biology easily become entangled with political and sociological problems—one does not have to be a fanatic to realize that the setting of research agendas and the parcelling out of resources involves the exercise of political power. All of us who pursued interests at the interface of physics and biology before it became popular have some personal perspectives on these issues. I will try to avoid these political entanglements and focus on our intellectual goals.

Looking back

Some of the giants of classical physics—Helmholtz, Maxwell, and Rayleigh, to name a few—routinely crossed borders among disciplines that we now distinguish as physics, chemistry, biology, and even psychology. Some of their forays into the phenomena of life were driven by a desire to test the universality of physical laws, such as the conservation of energy. A very different motivation was that our own view of the world is determined by what we can see and hear, and more subtly by what we can reliably infer from the data that our sense organs collect. These physicists thus were drawn to the study of the senses; for them, there was no boundary between optics and vision, or between acoustics and hearing. Helmholtz in particular took a very broad view, seeing a path not just from acoustics to the mechanics of the inner ear and from the properties of light to the optics of the eye, but all the way from the physical stimuli reaching our sense organs to the nature of our perceptions, to our ability to learn about the world, and even to what makes some sights or sounds more pleasing than others. Reading Helmholtz today I find myself struck by how much his insights still guide our thinking about vision and hearing, and by how the naturalness of his cross-disciplinary discourse remains something which few modern scientists achieve, despite all the current fanfare about the importance of multidisciplinary work. Most of all, I am struck by his soaring ambition that physics should not stop at the point where light hits our eyes or sound enters our ears, and that we should search for a physics that reaches all the way to our personal, conscious experience of the world in all its beauty.

The rise of modern physics motivated another wave of physicists to explore the phenomena of life. Fresh from the triumphs of quantum mechanics, they were emboldened to seek new challenges and brought new concepts. Bohr wondered aloud if the ideas of complementarity and indeterminacy would limit our ability to understand the microscopic events that provide the underpinnings of life. Delbrück was searching explicitly for new principles, hoping that a modern understanding of life would be as different from what came before as quantum mechanics was different from classical mechanics. Schrödinger, in his influential series of lectures entitled *What is Life?*, seized upon the discovery that our precious genetic inheritance was stored in objects the size of single molecules, highlighting how surprising this is for a classical physicist, and contrasted the order and complexity of life with the ordering of crystals, outlining a strikingly modern view of how non-equilibrium systems can generate structure out of disorder, continuously dissipating energy.

In one view of history, there is a direct path from Bohr, Delbrück and Schrödinger to the emergence of molecular biology. Certainly Delbrück did play a central role, not least because of his insistence that the community should

focus (as the physics tradition teaches us) on the simplest examples of crucial biological phenomena, reproduction and the transmission of genetic information. The goal of molecular biology to reduce these phenomena to interactions among a countable set of molecules surely echoed the physicists' search for the fundamental constituents of matter, and perhaps the greatest success of molecular biology is the discovery that many of these basic molecules of life are universal, shared across organisms separated by hundreds of millions of years of evolutionary history. Where classical biology emphasized the complexity and diversity of life, the first generation of molecular biologists emphasized the simplicity and universality of life's basic mechanisms, and it is not hard to see this as an influence of the physicists who came into the field at its start.

Another important idea at the start of molecular biology was that the structure of biological molecules matters. Although modern biology students, even in many high schools, can recite 'structure determines function,' this was not always obvious. To imagine, in the years immediately after World War II, that all of classical biochemistry and genetics would be reconceptualized once we could see the actual structures of proteins and DNA, was a revolutionary vision—a vision shared only by a handful of physicists and the most physical of chemists. Every physicist who visits the grand old Cavendish Laboratory in Cambridge should pause in the courtyard and realize that on that ground stood the 'MRC hut,' where Bragg nurtured a small group of young scientists who were trying to determine the structure of biological molecules through a combination of X-ray diffraction experiments and pure theory. To make a long and glorious story short, they succeeded, perhaps even beyond Bragg's wildest dreams, and some of the most important papers of twentieth century biology thus were written in a physics department.

Perhaps inspired by the successes of their intellectual ancestors, each subsequent generation of physicists offered a few converts. The idea, for example, that the flow of information through the nervous system might be reducible to the behavior of ion channels and receptors inspired one group, armed with low noise amplifiers, intuition about the interactions of charges with protein structure, and the theoretical tools to translate this intuition into testable, quantitative predictions. The possibility of isolating a single complex of molecules that carried out the basic functions of photosynthesis brought another group, armed with the full battery of modern spectroscopic methods that had emerged in solid state physics. Understanding that the mechanical forces generated by a focused laser beam are on the same scale as the forces generated by individual biological molecules as they go about their business brought another generation of physicists to our subject. The sequencing of whole genomes, including our own, generated the sense that

the phenomena of life could, at last, be explored comprehensively, and this inspired yet another group. These examples are far from complete, but give some sense for the diversity of challenges that drew physicists toward problems that traditionally had been purely in the domain of biologists.

Through these many generations, some conventional views arose about the nature of science at the borders between physics and biology. First, there is a strong emphasis on technique. From X-ray diffraction to the manipulation of single molecules to functional imaging of the brain, it certainly is true that physics has developed experimental techniques that allow much more direct exploration of questions raised by biologists. Second, there is a sense that in some larger classification system, biophysics is a biological science. Certainly when I was a student, and for many years afterwards, physicists would speak (sometimes wistfully) of colleagues who were fascinated by the phenomena of life as having “become biologists.” For their part, biologists would explain that physicists were successful in these explorations only to the extent that they appreciated what was “biologically important.” Finally, biophysics has come to be organized along the lines of the traditional biological subfields. As a result, the biophysics of neurons and the statistical mechanics of neural networks are separate subjects, and the generation of physicists exploring noise in the regulation of gene expression is disconnected from the previous generation that studied noise in ion channels.

Without taking anything away from what has been accomplished, I believe that much has been lost in the emergence of the conventional views about the nature of the interaction between physics and biology. By focusing on methods, we miss the fact that, faced with the same phenomena, physicists and biologists will ask different questions. In speaking of biological importance, we ignore the fact that physicists and biologists have different definitions of understanding. By organizing ourselves around structures that come from the history of biology, we lose contact with the dreams of our intellectual ancestors that the dramatic qualitative phenomena of life should be clues to deep theoretical insights, that there should be a physics of life and not just the physics of this or that particular process. It is, above all, these dreams that I would like to rekindle in my students and in the readers of this book.

Looking forward

At present, most questions about how things work in biological systems are viewed as questions that must be answered by experimental discovery. The situation in physics is very different, in that theory and experiment are more equal partners. In each area of physics we have a set of general theoretical principles, all interconnected,

which define what is possible; the path to confidence in any of these principles is built on a series of beautiful, quantitative experiments that have extended the envelope of what we can measure and know about the world. Beyond providing explanations for what has been seen, these principles provide a framework for exploring, sometimes playfully, what *ought* to be seen. In many cases these predictions are sufficiently startling that to observe the predicted phenomena (a new particle, a new phase of matter, fluctuations in the radiation left over from the big bang, ...) still constitutes a dramatic experimental discovery.

Can we imagine a physics of biological systems that reaches the level of predictive power that has become the standard in other areas of physics? Can we reconcile the physicists’ desire for unifying theoretical principles with the obvious diversity of life’s mechanisms? Could such theories engage meaningfully with the myriad experimental details of particular systems, yet still be derivable from succinct and abstract principles that transcend these details? For me, the answer to all of these questions is an enthusiastic “yes,” and I hope that this book will succeed in conveying both my enthusiasm and the reasons that lie behind it.

I have emphasized that, in the physics tradition, our subject should be defined by the kinds of questions we ask, but I haven’t given you a list of these questions. Worse yet, this emphasis on questions and concepts might leave us floating, disconnected from the data. It is, after all, the phenomena of life which are so dramatic and which demand our attention, so we should start there. There are so many beautiful things about life, however, that it can be difficult to choose a concrete starting point. Before explaining the choices I made in writing this book, I want to emphasize that there are many equally good choices. Indeed, if we choose almost any of life’s phenomena—the development of an embryo, our appreciation of music, the ability of bacteria to live in diverse environments, the way that ants find their way home in the hot desert—we can see glimpses of fundamental questions even in the seemingly most mundane events.

It is a remarkable thing that, pulling on the threads of one biological phenomenon, we can unravel so many general physics questions. In any one case, some problems will be presented in purer form than others, but in many ways everything is there. Thus, if we think hard about how crabs digest their food (to choose a particularly prosaic example), we will find ourselves worrying about how biological systems manage to find the right operating point in very large parameter spaces. This problem, as we will see in Chapter Three, arises in many different systems, across levels of organization from single protein molecules to short-term memory in the brain. Thus, in an odd way, everything is fair game. The challenge is not to find the most important or “fundamental”

phenomenon, but rather to see through any one of many interesting and beautiful phenomena to the deep physics problems that are hiding underneath the often formidable complexity of these systems.

The first problem, as noted above, is that there really is something different about being alive, and we'd like to know what this is—in the same way that we know what it is for a collection of atoms to be solid, for a collection of electrons to be superconducting, or for the vacuum to be confining (of quarks). This “What is life?” question harkens back to Schrödinger, and one might think that the molecular biology which arose in the decades after his manifesto would have answered his question, but this isn't clear. Looking around, we more or less immediately identify things which are alive, and the criteria that we use in making this discrimination between animate and inanimate matter surely have nothing to do with DNA or proteins. Even more strongly, we notice that things are alive long before we see them reproduce, so although self-reproduction might seem like a defining characteristic, it doesn't seem essential to our recognition of the living state. Being alive is a macroscopic state, while things like DNA and the machinery of self-reproduction are components of the microscopic mechanism by which this state is generated and maintained.⁴ While we have made much progress on identifying microscopic mechanisms, we have made rather less progress on identifying the “order parameters” that are characteristic of the macroscopic state.

Asking for the order parameters of the living state is a hard problem, and not terribly well posed. One way to make progress is to realize that as we make more quantitative models of particular biological systems, these models belong to families: we can imagine a whole class of systems, with varying parameters, of which the one we are studying is just one example. Presumably, most of these possible systems are not functional, living things. What then is special about the regions of parameter space that describe real biological systems? This is a more manageable question, and can be asked at many different levels of biological organization. If there is a principle that differentiates the genuinely biological parts of parameter space from the rest, then we can elevate this principle to a theory from which the properties of the biological system could be calculated a priori, as we do in other areas

⁴ More precisely, all the molecular components of life that we know about comprise *one way* of generating and maintaining the state that we recognize as being alive. We don't know if there are other ways, perhaps realized on other planets. This remark might once have seemed like science fiction, and perhaps it still is, but the discovery of planets orbiting distant stars has led many people to take these issues much more seriously. Designing a search for life on other planets gives us an opportunity to think more carefully about what it means to be alive.

of physics.

If real biological systems occupy only a small region in the space of possible systems, we have to understand the dynamics by which systems find their way to these special parameters. At one extreme, this is the problem of the origin of life. At the opposite extreme, we have the phenomena of physiological adaptation, whereby cells and systems adjust their behavior in relation to varying conditions or demands from the environment, sometimes in fractions of a second. In between we have learning and evolution. Adaptation, learning and evolution represent very different mechanisms, on different but perhaps overlapping time scales, for accomplishing a common goal, tuning the parameters of a biological system to match the problems that organisms need to solve as they try to survive and reproduce. What is the character of these dynamics? Are the systems that we see around us more or less “equilibrated” in these dynamics, or are today's organisms strongly constrained by the nature of the dynamics itself? Put another way, if evolution is implementing an algorithm for finding better organisms, are the functional behaviors of modern biological systems significantly shaped by the algorithm itself, or can we say that the algorithm solves a well defined problem, and what we see in life are the solutions to this problem?

In order to survive in the world, organisms do indeed have to solve a wide variety of problems. Many of these are really physics problems: converting energy from one form to another, sensing weak signals from the environment, controlling complex dynamical systems, transmitting information reliably from one place to another, or across generations, controlling the rates of thermally activated processes, predicting the trajectory of multidimensional signals, and so on. While it's obvious (now!) that everything which happens in living systems is constrained by the laws of physics, these physics problems in the life of the organism highlight these constraints and provide a special path for physics to inform our thinking about the phenomena of life.

Identifying all the physics problems that organisms need to solve is not so easy. Thinking about how single celled organisms, with sizes on the scale of one micron, manage to move through water, we quickly get to problems that have the look and feel of problems that we might find in Landau and Lifshitz. On the other hand, it really was a remarkable discovery that all cells have built Maxwell demons, and that our description of a wide variety of biochemical processes can be unified by this observation (see Section II.D). Efforts in this direction can be very rewarding, however, because we identify questions that connect functionally important behaviors—things organisms really care about, and for which evolution might select—with basic physical principles. Physics shows us what is hard about these problems, and where organisms face real challenges. In some cases, physics also places limits on what is possible, and

this gives us an opportunity to put the performance of biological systems on an absolute scale. This makes precise our intuition that organisms are really very good at solving some very difficult problems.

[I would like this paragraph to be better, but will come back to this.] To summarize, the business of life involves solving physics problems, and these problems provide us with a natural subject matter. In particular, these problems focus our attention on the concept of “function,” which is not part of the conventional physics vocabulary,⁵ but clearly is essential if we want to speak meaningfully about life. Of the possible mechanisms for solving these problems, most combinations of the available ingredients probably don’t work, and specifying this functional ensemble provides a manageable approach to the larger question of what characterizes the living state. Adaptation, learning and evolution allow organisms to find these special regions of parameter space, and the dynamics of these processes provide another natural set of problems.

If you are excited about problems at the interface of physics and biology, you must read Schrödinger’s “little book” *What is Life?*. To get a sense of the excitement and spirit of adventure that our intellectual ancestors brought to the subject, you should also look at the remarkable essays by Bohr (1933) and Delbrück (1949). Delbrück reflected on those early ideas some years later (1970), as did his colleagues and collaborators (Cairns et al 1966). For a more professional history of the emergence of modern molecular biology from these physicists’ musings, see Judson (1979).

Bohr 1933: Light and life. N Bohr, *Nature* **131**, 421–423 (1933).

Cairns et al 1966: *Phage and the Origins of Molecular Biology*, J Cairns, GS Stent & JD Watson, eds (Cold Spring Harbor Press, Cold Spring Harbor NY, 1966).

Delbrück 1949: A physicist looks at biology. M Delbrück, *Trans Conn Acad Arts Sci* **38**, 173–190 (1949). Reprinted in Cairns et al (1966), pp 9–22.

Delbrück 1970: A physicist’s renewed look at biology: twenty years later. M Delbrück, *Science* **168**, 1312–1315 (1970).

Judson 1979: *The Eighth Day of Creation* HF Judson (Simon and Schuster, New York, 1979).

Schrödinger 1944: *What is Life?* E Schrödinger (Cambridge University Press, Cambridge, 1944).

⁵ This isn’t quite fair. In thermodynamics we distinguish “useful work,” provides a notion of function, at least in the limited context of heat engines. But we need something much more general if we want to capture the full range of problems that organisms have to solve.

B. About this book

This book has its origins in a course that I have taught for several years at Princeton. It is aimed at PhD students in Physics, although a sizable number of brave undergraduates have also taken the course, as well as a handful of graduate students from biology, engineering, applied math, etc.. Bits and pieces have been tested in shorter courses, sometimes for quite different audiences, at the Marine Biological Laboratory, at Les Houches, at the Boulder Summer School on Condensed Matter Physics, at “Sapienza” Università di Roma, and at the Rockefeller University.

In early incarnations, the course consisted of a series of case studies—problems where physicists have tried to think about some particular biological system. The hope was that in each case study we might catch a glimpse of some deeper and more general ideas. As the course evolved, I tried to shift the balance from examples toward principles. The difficulty, of course, is that we don’t know the principles, we just have candidates. At some point I decided that this was OK, and that trying to articulate the principles was important even if we get them wrong. I believe that, almost by definition, something we will recognize as a theoretical physics of biological systems will have to cut across the standard subfields of biology, organizing our understanding of very different systems as instantiations of the same underlying ideas.

Although we are searching for principles, we start by being fascinated with the *phenomena* of life. Thus, the course starts with one particular biological phenomenon that holds, I think, an obvious appeal for physicists, and this is the ability of the visual system to count single photons. As we explore this phenomenon, we’ll meet some important facts about biological systems, we’ll see some methods and concepts that have wide application, and we’ll identify and sharpen a series of questions that we can recognize as physics problems. The really beautiful measurements that people have made in this system also provide a compelling antidote to the physicists’ prejudice that experiments on biological systems are necessarily messy; indeed, I think these measurements set a standard for quantitative experiments on biological systems that should be more widely appreciated and emulated.⁶

⁶ Perhaps surprisingly, many biologists share the expectation that their measurements will be noisy. Indeed, some biologists insist that physicists have to get used to this, and that this is a fundamental difference between physics and biology. Certainly it is a difference between the sciences as they are practiced, but the claim that there is something essentially sloppy about life is deeper, and deserves more scrutiny. One not so hidden agenda in my course is to teach physics students that it is possible to uncover precise, quantitative facts about biological systems in the same way that we can uncover precise quantitative facts about non-biological systems, and that this precision matters.

Another crucial feature of the photon counting problem is that it cuts across almost all levels of biological organization, from the quantum dynamics of single molecules to the macroscopic dynamics of human cognition.

Having introduced ourselves in some detail to one particular biological phenomenon, we proceed to explore three candidate principles: the importance of noise, the need for living systems to function without fine tuning of parameters, and the possibility that many of the different problems solved by living organisms are just different aspects of one big problem about the representation of information. Each of these ideas is something which many people have explored, and I hope to make clear that these ideas have generated real successes. The greatest successes, however, have been when these theoretical discussions are grounded in experiments on particular biological systems. As a result, the literature is fragmented along lines defined by the historical subfields of biology. The goal here is to present the discussion in the physics style, organized around principles from which we can derive predictions for particular examples.

My choice of candidate principles is personal, and I don't expect that everyone in the field will agree with me (see above). More importantly, the choice of examples is not meant to be canonical, but illustrative. In choosing these examples, I had three criteria. First, I had to understand what was going on, and of course this biases me toward cases which my friends and I have studied in the past. I apologize for this limitation, and hope that I have been able to do justice at least to some fraction of the field. Second, I want to emphasize the tremendous range of physics ideas which are relevant in thinking about the phenomena of life. Many students are given the impression, implicitly or explicitly, that to do biophysics one can get away with knowing less 'real physics' than in other subfields, and I think this is a disastrous misconception. Finally, if the whole program of finding principles is going to work, then it must be that a single principle really does illuminate the functioning of seemingly very different biological systems. Thus I make a special effort to be sure that the set of examples for each principle cuts across the subfields of biology, in particular across the great divide between molecular and cellular biology on the one hand and neurobiology on the other.

In trying to provide some perspective on our subject, in the previous section, I mentioned a number of now classic topics from across more than a century of interaction between physics and biology. I don't think it's right to teach by visiting these topics one after the other, for reasons which I hope are clear by now. On the other hand, it would be weird to take a whole course on biophysics and come out without having learned about these things. So I have tried to weave some of the classics into the conceptual framework of the course, perhaps sometimes in surprising places. There also are many beautiful things which I have left out, and again I apologize to peo-

ple who will find that I neglected matters close to their hearts. Sometimes the neglect reflects nothing more than my ignorance, but in some cases it is more subtle. I felt strongly that everything I discuss should fit together into a larger picture, and that it is almost disrespectful to give a laundry list of wonderful but undigested results. Thus, much was left unsaid.

I assume that readers (as with my students) have a strong physics background, and are comfortable with the associated mathematical tools. While many different areas of physics make an appearance, the most frequent references are to ideas from statistical mechanics. In practice, this is the area where at least U.S. students have the largest variance in their preparation. As a result, in places where my experience suggests that students will need help, I have not been shy to include (perhaps idiosyncratic) expositions of relevant physics topics that are not especially restricted to the biophysical context, since this is, after all, a physics course. Some more technical asides are presented as appendices. Throughout the text, and especially in the appendices, I try very hard to avoid saying "it can be shown that;" the resulting text is longer, but I hope more useful.

No matter how much we may be searching for deep theoretical principles, in the physics tradition, we do need a grasp of the facts. But when we teach particle physics we don't start by reading from the particle data book, so similarly I don't start by reciting the "biological background." Rather, we plunge right in, and as we encounter things that need explaining, I try to explain them. I do want to emphasize (maybe this is especially meaningful coming from a theorist!) the importance of mastering the experimental facts about systems that we find interesting. I think we should avoid talking about how "physicists need to learn the biology," since "biology" could mean either the study of living systems or the academic discipline practiced in biology departments, and these need not be the same thing. We must know what has been measured, assess these data with informed skepticism, and use the results to guide our thinking as we ask our own new and interesting questions. I hope I manage to strike the right balance.

The most important comment about the structure of the book concerns the problems. I cannot overstate the importance of doing problems as a component of learning. One should go further, getting into the habit of calculating as one reads, checking that you understand all the steps of an argument and that things make sense when you plug in the numbers or make order of magnitude estimates. For all these reasons, I have chosen (following Landau and Lifshitz) to embed the problems in the text, rather than relegating them to the ends of chapters. In some places the problems are small, really just reminding you to fill in some missing steps before going on to the next topic. At the opposite extreme, some problems are small research projects. Be-

cause progress in biophysics depends on intimate interaction between theory and experiment, some of the problems ask you to analyze real data, which can be found at <http://www.princeton.edu/~wbialek/PHY562/data>.

Let me also say a few words about references. References to the original literature serve multiple functions, especially in textbooks. Most obviously, I should cite the papers that most influenced my own thinking about the subject, acknowledging my intellectual debts. Since this text is based on a course for PhD students, citations also help launch the student into the current literature, marking the start of paths that can carry you well beyond digested, textbook discussions. In another direction, references point back to classic papers, papers worth reading decades after they were published, papers that can provide inspiration. Importantly, all of these constitute subjective criteria for inclusion on the reference list, and so I think it is appropriate to collect references with some commentary, as you have already seen at the end of the previous section. Let me note that the reference list should *not* be viewed as a rough draft of the history of the subject, nor as an attempt to establish objective priorities for some work over others.

C. About this draft

This is not the final draft of the book. I know there are things that need to be fixed, but I have been pushing to get the text to the point where I won't be embarrassed by letting other people look at it (I hope!). My own concerns about the state of the text include the following:

1. All the figures are placeholders. Some are grabbed from published papers, while others are bad photographs of what I sketched on the blackboard. There is work to be done in bringing all of this up to a standard of clarity and consistency.
2. I have pushed through the text several times, but I haven't really had a chance to look at the balance of topics. I worry that things which I know best have grown out of proportion to other topics, and I could use some advice. There is a related question about which things belong in the main text and which can be safely pushed to the Appendices.
3. There are places where I just haven't finished, even if I am pretty sure what needs to be done. This has been a very long project, but I fully expect readers to give advice that will necessitate further revision. Thus, I thought it might be OK to let people see things even with the gaps—perhaps you even have ideas about how to fill them in. These problem areas of the text are flagged in **red**. In some places these are small (I think) nagging questions, while in other areas there are bigger sections missing.

4. I have been working hard on the opening parts of chapters and sections, trying to provide more context and a guide to what is coming. The ends of many sections still seem a bit abrupt, however, suggesting that I might have stopped when I was exhausted by the topic rather than when I reached a conclusion. This will get fixed.

At this stage of the project, all input is welcome. I hope you will read sympathetically as well as critically, but getting things right is important, so feel free to bash away.

ACKNOWLEDGMENTS

Even if I had the perfect idea for teaching a course, it would be meaningless without students. By now, hundreds of students have listened to the whole set of lectures and worked through the problems, providing feedback at every stage, as have several very able teaching assistants. At least as many students have heard pieces of the course, in different venues, and every time I taught I learned something—at least, I hope, about how to say things more clearly. Less tangible, but even more important, the liveliness and engagement of the students have made teaching a pleasure.

The views of the field which I present here are personal, and I don't want anyone else held responsible for my foibles. On the other hand, these views did not emerge in isolation. I am especially grateful to Rob de Ruyter van Steveninck, who introduced me to the wonders of close collaboration between theory and experiment. What began as a brief discussion about the possibility of measuring the precision of computation in a small corner of the fly's brain has become half a lifetime of friendship and shared intellectual adventure.

My good fortune in finding wonderful experimental collaborators began with Rob, but certainly didn't end there. A decade of conversations with Michael Berry, Allison Doupe, Steve Lisberger and Leslie Osborne, sometimes reflected in joint papers and sometimes not, have all influenced important parts of this book, in ways which I hope they will recognize. After I moved to Princeton, David Tank, Eric Wieschaus and I began a very different adventure, soon joined by Thomas Gregor. I have been amazed by how these interactions have so quickly reshaped my own thinking, leaving their mark on my view of the subject as a whole and hence on this text.

Theory itself is more fun in collaboration with others, even when we aren't engaged with our experimental friends. Different parts of the text trace their origins to joint work with N Brenner, WJ Bruno, CG Callan, M DeWeese, AL Fairhall, S Kivelson, R Koberle, T Mora, I Nemenman, JN Onuchic, SE Palmer, M Potters, FM Rieke, DL Ruderman, E Schneidman, S Setayeshgar, T Sharpee, GJ Stephens, S Still, SP Strong, G Tkačik, N

Tishby, A Walczak, D Warland and A Zee. I am hugely grateful to all of them.

It is almost embarrassing to admit that I first taught PHY 562 a very long time ago, while I was still a member of the NEC Research Institute, and a visiting lecturer at Princeton. Dawon Kahng and Joe Giordmaine were responsible for creating the enlightened environment at NEC, which lasted for a marvelous decade, while David Gross and Stew Smith made it possible for me to teach those early versions of the course at Princeton. The opportunity to interact with students while still enjoying the support of an industrial research laboratory dedicated to basic science was quite magical. During this period, frequent discussions with Albert Libchaber were also important, as he insisted that explorations at the interface of physics and biology be ambitious but still crisp and decisive—a demanding combination.

Although the wonders of life in industrial labs have largely disappeared, the pleasures of teaching at Princeton have continued and grown. I am especially grateful to my colleagues in the Physics department for welcoming the intellectual challenges posed by the phenomena of life as being central to physics itself, rather than being “applications” of physics to another field. The result has been the coalescence of a very special community, and I hope that some of what I have learned from this community is recorded faithfully in this book. John Hopfield’s role in making all this happen—by setting an example for what could be done, by being an explicit (and horrifyingly witty) provocateur, and by being a quiet but persistent catalyst for change—cannot be overestimated; it a pleasure to thank him. I don’t think that even John imagined that there would eventually be a “biophysics theory group” at Princeton, but with Curt Callan and Ned Wingreen, we have managed to do it, and we have been joined by a succession of young colleagues who have held the Lewis–Sigler Fellowship—M Desai, J England and M Kaschube—all of whom have added enormously to our community. Curt deserves special thanks, for his leadership and even more for the energy and enthusiasm he brings to seminars and discussions, engaging with the details but also reminding us that theoretical physics has lofty aspirations.

Everyone who has tried to write a book based on their teaching experience knows the enormous difference between a good set of lecture notes and the final product. I very much appreciate Arthur Wightman’s suggestion, long ago, that this transition would be worth the effort. Ingrid Gnerlich, my editor at Princeton University Press, has consistently provided the right combination of encouragement and gentle reminders of looming (and passing) deadlines. The idea of actually finishing (!) started to crystallize during a wonderful sabbatical in Rome, and has been greatly helped along by visiting professorships at the Rockefeller University and most recently at The Graduate Center of the City University of New York.

Both in Rome and in New York, stimuli from colleagues and from the surrounding cities have proved delightfully synergistic.

Despite my reservations (see above), I am much more comfortable with this draft than I was with the previous one, and this is the result of wonderful input on short notice from several colleagues. Rob Phillips brought objectivity, and the proper amount of scathing humor, alerting me to a variety of problems. Thomas Gregor, Justin Kinney and Fred Rieke gave generously of their expertise, and Rob de Ruyter provided yet more of the insight, craftsmanship and knowledge of scientific history that I have so much enjoyed in our long collaboration. My thanks to all of them.

It often is remarked that theory is a relatively inexpensive activity, so that we theorists are less dependent on raising money than are our experimentalist friends. But theory is a communal activity, and all the members of the community need salaries. Because I have benefited so much from the stimulation provided by the scientists around me, I am especially grateful for the steady support my colleagues and I have received from the National Science Foundation, and for the generosity of Princeton University in bringing all of us together. In particular, Denise Caldwell, Kenneth Whang and especially Kras-tan Blagoev deserve our thanks for helping to insure that this kind of science has a home at the NSF, even in difficult times. The Burroughs–Wellcome Fund, the WM Keck Foundation, and the Swartz Foundation have also been extremely generous, sometimes leaping in where the usual angels feared to tread.

Finally, while the product of the scientific enterprise must have meaning outside our individual feelings, the process of science is intensely personal. When we collaborate or even just learn from one another, we share not just our ideas about the next step in a small project, but our hopes and dreams for efforts that could occupy a substantial fraction of a lifetime. To make progress we admit to one another how little we understand, and how we struggle even to formulate the questions. For want of a better word, collaboration is an intimate activity. Colleagues become friends, friendships deepen, we come to care not just about ideas and results but about one another. It is, by any measure, a privileged life. If this text helps some readers to find their way to such enjoyment, I will have repaid a small fraction of my debt.

William Bialek
September 18, 2011

I. PHOTON COUNTING IN VISION

Imagine sitting quietly in a dark room, staring straight ahead. A light flashes. Do you see it? Surely if the flash is bright enough the answer is yes, but how dim can the flash be before we fail? Do we fail abruptly, so that there is a well defined threshold—lights brighter than threshold are always seen, lights dimmer than threshold are never seen—or is the transition from seeing to not seeing somehow more gradual? These questions are classical examples of “psychophysics,” studies on the relationship between perception and the physical variables in the world around us, and have a history reaching back at least into the nineteenth century.

In 1911, the physicist Lorentz was sitting in a lecture that included an estimate of the “minimum visible,” the energy of the dimmest flash of light that could be consistently and reliably perceived by human observers. But by 1911 we knew that light was composed of photons, and if the light is of well defined frequency or wavelength then the energy E of the flash is equivalent to an easily calculable number of photons n , $n = E/h\nu$. Doing this calculation, Lorentz found that just visible flashes of light correspond to roughly 100 photons incident on our eyeball. Turning to his physiologist colleague Zwaa-dermaker, Lorentz asked if much of the light incident on the cornea might get lost (scattered or absorbed) on its way through the gooey interior of the eyeball, or if the experiments could be off by as much as a factor of ten. In other words, is it possible that the real limit to human vision is the counting of single photons?

Lorentz’ suggestion really is quite astonishing. If correct, it would mean that the boundaries of our perception are set by basic laws of physics, and that we reach the limits of what is possible. Further, if the visual system really is sensitive to individual light quanta, then some of the irreducible randomness of quantum events should be evident in our perceptions of the world around us, which is a startling thought.

In this Chapter, we will see that humans (and other animals) really can detect the arrival of individual photons at the retina. Tracing through the many steps from photon arrival to perception we will see a sampling of the physics problems posed by biological systems, ranging from the dynamics of single molecules through amplification and adaptation in biochemical reaction networks, coding and computation in neural networks, all the way to learning and cognition. For photon counting some of these problems are solved, but even in this well studied case many problems are open and ripe for new theoretical and experimental work. The problem of photon counting also introduces us to methods and concepts of much broader applicability. We begin by exploring the phenomenology, aiming at the formulation of the key physics problems. By the end of the Chapter I hope to have formulated an approach to the exploration of bio-

logical systems more generally, and identified some of the larger questions that will occupy us in Chapters to come.

A. Posing the problem

One of the fundamental features of quantum mechanics is randomness. If we watch a single molecule and ask if it absorbs a photon, this is a random process, with some probability per unit time related to the light intensity. The emission of photons is also random, so that a typical light source does not deliver a fixed number of photons. Thus, when we look at a flash of light, the number of photons that will be absorbed by the sensitive cells in our retina is a random number, not because biology is noisy but because of the physics governing the interaction of light and matter. One way of testing whether we can count single photons, then, is to see if we can detect the signatures of this quantum randomness in our perceptions. This line of reasoning came to fruition in experiments by Hecht, Schlaer and Pirenne (in New York) and by van der Velden (in the Netherlands) in the early 1940s. [Need to check what was done by Barnes & Czerny, between Lorentz and 1940s]

What we think of classically as the intensity of a beam of light is proportional to the *mean* number of photons per second that arrive at some region where they can be counted.⁷ For most conventional light sources, however, the stream of photons is not regular, but completely random. Thus, in any very small window of time dt , there is a probability $r dt$ that a photon will be counted, where r is the mean counting rate or light intensity, and the events in different time windows are independent of one another. These are the defining characteristics of a “Poisson process,” which is the maximally random sequence of point events in time—if we think of the times at which photons are counted as being like the positions of particles, then the sequence of photon counts is like an ideal gas, with no correlations or “interactions” among the particles at different positions.

As explained in detail in Appendix A.1, if events occur as a Poisson process with rate r , then if we count the events over some time T , the mean number of counts will be $M = rT$, but the probability that we actually obtain a count of n will be given by the Poisson distribution,

$$P(n|M) = e^{-M} \frac{M^n}{n!}. \quad (1)$$

In our case, the mean number of photons that will be counted at the retina is proportional to the classical intensity of the light flash, $M = \alpha I$, where the constant

⁷ More precisely, we can measure the mean number of photons per second per unit area.




FIG. 1 Probability of seeing calculated from Eq. (2), where the intensity I is measured as the mean number of photons incident on the cornea, so that α is dimensionless. Curves are shown for different values of the threshold photon count K and the scaling factor α . Note the distinct shapes for different K , but when we change α at fixed K we just translate the curve along the the log intensity axis, as shown by the red dashed arrow.

α includes all the messy details of what happens to the light on its way through the eyeball.⁸ Thus, when we deliver the “same” flash of light again and again, the actual physical stimulus to the retina will fluctuate, and it is plausible that our perceptions will fluctuate as well.

Let’s be a bit more precise about all of this. In the simplest view, you would be willing to say “yes, I saw the flash” once you had counted K photons. Equation (1) tell us the probability of counting exactly n photons given the mean, and the mean is connected to the intensity of the flash by $M = \alpha I$. Thus we predict that there is a probability of seeing a flash of intensity I ,

$$P_{\text{see}}(I) = \sum_{n=K}^{\infty} P(n|M = \alpha I) = e^{-\alpha I} \sum_{n=K}^{\infty} \frac{(\alpha I)^n}{n!}. \quad (2)$$

So, if we sit in a dark room and watch as dim lights are flashed, we expect that our perceptions will fluctuate—sometimes we see the flash and sometimes we don’t—but there will be an orderly dependence of the *probability* of seeing on the intensity, given by Eq (2). Importantly, if

⁸ The units for light intensity are especially problematic. Today we know that we should measure the number of photons arriving per second, per unit area, but many of the units were set before this was understood. Also, if we have a broad spectrum of wavelengths, we might want to weight the contributions from different wavelengths not just by their contribution to the total energy but by their contribution to the overall appearance of brightness. Thus, some of the complications have honest origins.

we plot P_{see} vs. $\log I$, as in Fig. 1, then the *shape* of the curve depends crucially on the threshold photon count K , but changing the unknown constant α just translates the curve along the x-axis. So we have a chance to measure the threshold K by looking at the shape of the curve; more fundamentally we might say we are testing the hypothesis that the probabilistic nature of our perceptions is determined by the physics of photon counting.

Problem 1: Photon statistics, part one. There are two reasons why the arrival of photons might be described by a Poisson process. The first is a very general “law of small numbers” argument. Imagine a general point process in which events occur at times $\{t_i\}$, with some correlations among the events. Assume that these correlations die away with some correlation time, so that events i and j are independent if $|t_i - t_j| \gg \tau_c$. Explain qualitatively why, if we select events out of the original sequence at random, then if we select a sufficiently small fraction of these events the resulting sparse sequence will be approximately Poisson. What is the condition for the Poisson approximation to be a good one? What does this have to do with why, for example, the light which reaches us from an incandescent light bulb comes in a Poisson stream of photons?

Problem 2: How many sources of randomness? As noted above, the defining feature of a Poisson process is the independence of events at different times, and typical light sources generate a stream of photons whose arrival times approximate a Poisson process. But when we count these photons, we don’t catch every one. Show that if the photon arrivals are a Poisson process with rate r , and we count a fraction f these, selected at random, then the times at which events are counted will also be a Poisson process, with rate fr . Why doesn’t the random selection of events to be counted result in some “extra” variance beyond expectations for the Poisson process?

Hecht, Shlaer and Pirenne did exactly the experiment we are analyzing. Subjects (the three co-authors) sat in a dark room, and reported whether they did or did not see a dim flash of light. For each setting of the intensity, there were many trials, and responses were variable, but the subjects were forced to say yes or no, with no “maybe.” Thus, it was possible to measure at each intensity the probability that the subject would say yes, and this is plotted in Fig 2.

The first nontrivial result of these experiments is that human perception of dim light flashes really is probabilistic. No matter how hard we try, there is a range of light intensities in which our perceptions fluctuate from flash to flash of the same intensity, seeing one and missing another. Quantitatively, the plot of probability of seeing vs $\log(\text{intensity})$ is fit very well by the predictions from the Poisson statistics of photon arrivals. In particular, Hecht, Shlaer and Pirenne found a beautiful fit in the range from $K = 5$ to $K = 7$; subjects of different age had very different values for α (as must be true if light transmission through the eye gets worse with age) but




FIG. 2 Probability of seeing calculated from Eq. (2), with the threshold photon count $K = 6$, compared with experimental results from Hecht, Shlaer and Pirenne. For each observer we can find the value of α that provides the best fit, and then plot all the data on a common scale as shown here. Error bars are computed on the assumption that each trial is independent, which probably generates errors bars that are slightly too small.

similar values of K . In Fig 2 I've shown all three observers' data fit to $K = 6$, along with error bars (absent in the original paper); although one could do better by allowing each person to have a different value of K , it's not clear that this would be supported by the statistics. The different values of α , however, are quite important.

Details aside, the frequency of seeing experiment brings forward a beautiful idea: the probabilistic nature of our perceptions just reflects the physics of random photon arrivals. An absolutely crucial point is that Hecht, Shlaer and Pirenne chose stimulus conditions such that the 5 to 7 photons needed for seeing are distributed across a broad area on the retina, an area that contains hundreds of photoreceptor cells [perhaps this needs to be clearer?] Thus the probability of one receptor (rod) cell getting more than one photon is very small. The experiments on human behavior therefore indicate that individual photoreceptor cells generate reliable responses to single photons. In fact, vision begins (as we discuss in more detail soon) with the absorption of light by the visual pigment rhodopsin, and so sensitivity to single photons means that each cell is capable of responding to a single molecular event. This is a wonderful example of using macroscopic experiments to draw conclusions about single cells and their microscopic mechanisms.

Problem 3: Simulating a Poisson process. Much of what we want to know about Poisson processes can be determined analytically (see Appendix A.1). Thus if we do simulations we know

what answer we should get (!). This provides us with an opportunity to exercise our skills, even if we don't get any new answers. In particular, *doing* a simulation is never enough; you have to analyze the results, just as you analyze the results of an experiment. Now is as good a time as any to get started. If you are comfortable doing everything in a programming language like C or Fortran, that's great. On the other hand, high-level languages such as MATLAB or Mathematica have certain advantages. Here you should use MATLAB to simulate a Poisson process, and then analyze the results to be sure that you actually did what you expected to do. [Before finalizing, check on the use of free version of MATLAB, Octave.]

(a) MATLAB has a command `rand` that generates random numbers with a uniform distribution from 0 to 1. Consider a time window of length T , and divide this window into many small bins of size dt . In each bin you can use `rand` to generate a number which you can compare with a threshold—if the random number is above threshold you put an event in the bin, and you can adjust the threshold to set the average number of events in the window. You might choose $T = 10^3$ sec and arrange that the average rate of the events is $\bar{r} \sim 10$ per second; note that you should be able to relate the threshold to the mean rate \bar{r} analytically. Notice that this implements (in the limit $dt \rightarrow 0$) the definition of the Poisson process as independent point events.

(b) The next step is to check that the events you have made really do obey Poisson statistics. Start by counting events in windows of some size τ . What is the mean count? The variance? Do you have enough data to fill in the whole probability distribution $P_\tau(n)$ for counting n of events in the window? How do all of these things change as you change τ ? What if you go back and make events with a different average rate? Do your numerical results agree with the theoretical expressions? In answering this question, you could try to generate sufficiently large data sets that the agreement between theory and experiment is almost perfect, but you could also make smaller data sets and ask if the agreement is good within some estimated error bars; this will force you to think about how to put error bars on a probability distribution. [Do we need to have some more about error bars somewhere in the text?] You should also make a histogram (`hist` should help) of the times between successive events; this should be an exponential function, and you should work to get this into a form where it is a properly normalized probability density. Relate the mean rate of the events to the shape of this distribution, and check this in your data.

(c) Instead of deciding within each bin about the presence or absence of an event, use the command `rand` to choose N random times in the big window T . Examine as before the statistics of counts in windows of size $\tau \ll T$. Do you still have an approximately Poisson process? Why? Do you see connections to the statistical mechanics of ideal gases and the equivalence of ensembles?

Problem 4: Photon statistics, part two. The other reason why we might find photon arrivals to be a Poisson process comes from a very specific quantum mechanical argument about coherent states. This argument may be familiar from your quantum mechanics courses, but this is a good time to review. If you are not familiar with the description of the harmonic oscillator in terms of raising and lowering or creation and annihilation operators, try the next problem, which derives many of the same conclusions by making explicit use of wave functions.

(a.) We recall that modes of the electromagnetic field (in a free space, in a cavity, or in a laser) are described by harmonic oscillators. The Hamiltonian of a harmonic oscillator with frequency ω can be written as

$$\mathbf{H} = \hbar\omega(a^\dagger a + 1/2), \quad (3)$$

where a^\dagger and a are the creation and annihilation operators that connect states with different numbers of quanta,

$$a^\dagger|n\rangle = \sqrt{n+1}|n+1\rangle, \quad (4)$$

$$a|n\rangle = \sqrt{n}|n-1\rangle. \quad (5)$$

There is a special family of states called coherent states, defined as

eigenstates of the annihilation operator,

$$a|\alpha\rangle = \alpha|\alpha\rangle. \quad (6)$$

If we write the coherent state as a superposition of states with different numbers of quanta,

$$|\alpha\rangle = \sum_{n=0}^{\infty} \psi_n |n\rangle, \quad (7)$$

then you can use the defining Eq (6) to give a recursion relation for the ψ_n . Solve this, and show that the probability of counting n quanta in this state is given by the Poisson distribution, that is

$$P_\alpha(n) \equiv \left| \langle n|\alpha\rangle \right|^2 = |\psi_n|^2 = e^{-M} \frac{M^n}{n!}, \quad (8)$$

where the mean number of quanta is $M = |\alpha|^2$.

(b.) The specialness of the coherent states relates to their dynamics and to their representation in position space. For the dynamics, recall that any quantum mechanical state $|\phi\rangle$ evolves in time according to

$$i\hbar \frac{d|\phi\rangle}{dt} = \mathbf{H}|\phi\rangle. \quad (9)$$

Show that if the system starts in a coherent state $|\alpha(0)\rangle$ at time $t = 0$, it remains in a coherent state for all time. Find $\alpha(t)$.

(c.) If we go back to the mechanical realization of the harmonic oscillator as a mass m hanging from a spring, the Hamiltonian is

$$\mathbf{H} = \frac{1}{2m} p^2 + \frac{m\omega^2}{2} q^2, \quad (10)$$

where p and q are the momentum and position of the mass. Remind yourself of the relationship between the creation and annihilation operators and the position and momentum operators (\hat{q}, \hat{p}). In position space, the ground state is a Gaussian wave function,

$$\langle q|0\rangle = \frac{1}{(2\pi\sigma^2)^{1/4}} \exp\left(-\frac{q^2}{4\sigma^2}\right), \quad (11)$$

where the variance of the zero point motion $\sigma^2 = \hbar/(4m\omega)$. The ground state is also a “minimum uncertainty wave packet,” so called because the variance of position and the variance of momentum have a product that is the minimum value allowed by the uncertainty principle; show that this is true. Consider the state $|\psi(q_0)\rangle$ obtained by displacing the ground state to a position q_0 ,

$$|\psi(q_0)\rangle = e^{iq_0\hat{p}}|0\rangle. \quad (12)$$

Show that this is a minimum uncertainty wave packet, and also a coherent state. Find the relationship between the coherent state parameter α and the displacement q_0 .

(d.) Put all of these steps together to show that the coherent state is a minimum uncertainty wave packet with expected values of the position and momentum that follow the classical equations of motion.

Problem 5: Photon statistics, part two, with wave functions. Work out a problem that gives the essence of the above using wave functions, without referring to α and a^\dagger .

There is a very important point in the background of this discussion. By placing results from all three observers on the same plot, and fitting with the same value of K , we are claiming that there is something reproducible, from individual to individual, about our perceptions. On the other hand, the fact that each observer has a different value for α means that there are individual differences, even in this simplest of tasks. Happily, what seems to be reproducible is something that feels

like a fundamental property of the system, the number of photons we need to count in order to be sure that we saw something. But suppose that we just plot the probability of seeing vs the (raw) intensity of the light flash. If we average across individuals with different α s, we will obtain a result which does not correspond to the theory, and this failure might even lead us to believe that the visual system does not count single photons. This shows us that (a) finding what is reproducible can be difficult, and (b) averaging across an ensemble of individuals can be *qualitatively* misleading. Here we see these conclusions in the context of human behavior, but it seems likely that similar issues arise in the behavior of single cells. The difference is that techniques for monitoring the behavior of single cells (e.g., bacteria), as opposed to averages over populations of cells, have emerged much more recently. As an example, it still is almost impossible to monitor, in real time, the metabolism of single cells, whereas simultaneous measurements on many metabolic reactions, averaged over populations of cells, have become common. We still have much to learn from these older experiments!

Problem 6: Averaging over observers. Go back to the original paper by Hecht, Shlaer and Pirenne⁹ and use their data to plot, vs. the intensity of the light flash, the probability of seeing averaged over all three observers. Does this look anything like what you find for individual observers? Can you simulate this effect, say in a larger population of subjects, by assuming that the factor α is drawn from a distribution? Explore this a bit, and see how badly misled you could be. This is not too complicated, I hope, but deliberately open ended.

Before moving on, a few more remarks about the history. [I have some concern that this is a bit colloquial, and maybe more like notes to add to the references than substance for the text. Feedback is welcome here.] It's worth noting that van der Velden's seminal paper was published in Dutch, a reminder of a time when anglophone cultural hegemony was not yet complete. Also (maybe more relevant for us), it was published in a physics journal. The physics community in the Netherlands during this period had a very active interest in problems of noise, and van der Velden's work was in this tradition. In contrast, Hecht was a distinguished contributor to understanding vision but had worked within a “photochemical” view which he would soon abandon as inconsistent with the detectability of single photons and hence single molecules of activated rhodopsin. Parallel

⁹ As will be true throughout the text, references are found at the end of the section.




FIG. 3 (a) A single rod photoreceptor cell from a toad, in a suction pipette. Viewing is with infrared light, and the bright bar is a stimulus of 500 nm light. (b) Equivalent electrical circuit for recording the current across the cell membrane [really needs to be redrawn, with labels!]. (c) Mean current in response to light flashes of varying intensity. Smallest response is to flashes that deliver a mean ~ 4 photons, successive flashes are brighter by factors of 4. (d) Current responses to repeated dim light flashes at times indicated by the tick marks. Note the apparently distinct classes of responses to zero, one or two photons. From Rieke & Baylor (1998).

to this work, Rose and de Vries (independently) emphasized that noise due to the random arrival of photons at the retina also would limit the reliability of perception at intensities well above the point where things become barely visible. In particular, de Vries saw these issues as part of the larger problem of understanding the physical limits to biological function, and I think his perspective on the interaction of physics and biology was far ahead of its time.

It took many years before anyone could measure directly the responses of photoreceptors to single photons. It was done first in the (invertebrate) horseshoe crab [be sure to add refs to Fuortes & Yeandle; maybe show a figure?], and eventually by Baylor and coworkers in toads and then in monkeys. The complication in the lower vertebrate systems is that the cells are coupled together, so that the retina can do something like adjusting the size of pixels as a function of light intensity. This means that the nice big current generated by one cell is spread as a small voltage in many cells, so the usual method of measuring the voltage across the membrane of one cell won't work; you have to suck the cell into a pipette and collect the current, as seen in Fig 3.

Problem 7: Gigaseals. As we will see, the currents that are relevant in biological systems are on the order of picoAmps.

Although the response of rods to single photons is slow, many processes in the nervous system occur on the millisecond timescale. Show that if we want to resolve picoAmps in milliseconds, then the leakage resistance (e.g. between rod cell membrane and the pipette in Fig 3) must be $\sim 10^9$ ohm, to prevent the signal being lost in Johnson noise.

In complete darkness, there is a 'standing current' of roughly 20 pA flowing through the membrane of the rod cell's outer segment. You should keep in mind that currents in biological systems are carried not by electrons or holes, as in solids, but by ions moving through water; we will learn more about this below [be sure we do!]. In the rod cell, the standing current is carried largely by sodium ions, although there are contributions from other ions as well. This is a hint that the channels in the membrane that allow the ions to pass are not especially selective for one ion over the other. The current which flows across the membrane of course has to go somewhere, and in fact the circuit is completed within the rod cell itself, so that what flows across the outer segment of the cell is compensated by flow across the inner segment [improve the figures to show this clearly]. When the rod cell is exposed to light, the standing current is reduced, and with sufficiently bright flashes it is turned off all together.

As in any circuit, current flow generates changes in the voltage across the cell membrane. Near the bottom of the cell [should point to better schematic, one figure with everything we need for this paragraph] there are special channels that allow calcium ions to flow into the cell in response to these voltage changes, and calcium in turn triggers the fusion of vesicles with the cell membrane. These vesicles are filled with a small molecule, a neurotransmitter, which can then diffuse across a small cleft and bind to receptors on the surface of neighboring cells; these receptors then respond (in the simplest case) by opening channels in the membrane of this second cell, allowing currents to flow. In this way, currents and voltages in one cell are converted, via a chemical intermediate, into currents and voltages in the next cell, and so on through the nervous system. The place where two cells connect in this way is called a synapse, and in the retina the rod cells form synapses onto a class of cells called bipolar cells. More about this later, but for now you should keep in mind that the electrical signal we are looking at in the rod cell is the first in a sequence of electrical signals that ultimately are transmitted along the cells in the optic nerve, connecting the eye to the brain and hence providing the input for visual perception.

Very dim flashes of light seem to produce a quantized reduction in the standing current, and the magnitude of these current pulses is roughly 1 pA, as seen in Fig 3. When we look closely at the standing current, we see that it is fluctuating, so that there is a continuous background noise of ~ 0.1 pA, so the quantal events are




FIG. 4 A closer look at the currents in toad rods. At left, five instances in which the rod is exposed to a dim flash at $t = 0$. It looks as if two of these flashes delivered two photons (peak current ~ 2 pA), one delivered one photon (peak current ~ 1 pA), and two delivered zero. The top panel shows the raw current traces, and the bottom panel shows what happens when we smooth with a 100 ms window to remove some of the high frequency noise. At right, the distribution of smoothed currents at the moment t_{peak} when the average current peaks; the data (circles) are accumulated from 350 flashes in one cell, and the error bars indicate standard errors of the mean due to this finite sample size. Solid green line is the fit to Eq (19), composed of contributions from $n = 0, n = 1, \dots$ photon events, shown red. Dashed blue lines divide the range of observed currents into the most likely assignments to different photon counts. These data are from unpublished experiments by FM Rieke at the University of Washington; many thanks to Fred for providing the data in raw form.

easily detected. It takes a bit of work to convince yourself that these events really are the responses to single photons. Perhaps the most direct experiment is to measure the cross-section for generating the quantal events, and compare this with the absorption cross-section of the rod, showing that a little more than 2/3 of the photons which are absorbed produce current pulses. In response to steady dim light, we can observe a continuous stream of pulses, the rate of the pulses is proportional to the light intensity, and the intervals between pulses are distributed exponentially, as expected if they represent the responses to single photons (cf Section A.1).

Problem 8: Are they really single photon responses?

Work out a problem to ask what aspects of experiments in Fig 4 are the smoking gun. In particular, if one pulse were from the coincidence of two photons, how would the distribution of peak heights shift with changing flash intensity?

When you look at the currents flowing across the rod cell membrane, the statement that single photon events

are detectable above background noise seems pretty obvious, but it would be good to be careful about what we mean here. In Fig 4 we take a closer look at the currents flowing in response to dim flashes of light. These data were recorded with a very high bandwidth, so you can see a lot of high frequency noise. Nonetheless, in these five flashes, it's pretty clear that twice the cell counted zero photons, once it counted one photon (for a peak current ~ 1 pA) and twice it counted two photons; this becomes even clearer if we smooth the data to get rid of some of the noise. Still, these are anecdotes, and one would like to be more quantitative.

Even in the absence of light there are fluctuations in the current, and for simplicity let's imagine that this background noise is Gaussian with some variance σ_0^2 . The simplest way to decide whether we saw something is to look at the rod current at one moment in time, say at $t = t_{\text{peak}} \sim 2$ s after the flash, where on average the current is at its peak. Then given that no photons were counted, this current i should be drawn out of the probability distribution

$$P(i|n=0) = \frac{1}{\sqrt{2\pi\sigma_0^2}} \exp\left[-\frac{i^2}{2\sigma_0^2}\right]. \quad (13)$$

If one photon is counted, then there should be a mean current $\langle i \rangle = i_1$, but there is still some noise. Plausibly the noise has two pieces—first, the background noise still is present, with its variance σ_0^2 , and in addition the amplitude of the single photon response itself can fluctuate; we assume that these fluctuations are also Gaussian and independent of the background, so they just add σ_1^2 to the variance. Thus we expect that, in response to one photon, the current will be drawn from the distribution

$$P(i|n=1) = \frac{1}{\sqrt{2\pi(\sigma_0^2 + \sigma_1^2)}} \exp\left[-\frac{(i - i_1)^2}{2(\sigma_0^2 + \sigma_1^2)}\right]. \quad (14)$$

If each single photon event is independent of the others, then we can generalize this to get the distribution of currents expected in response of $n = 2$ photons, [need to explain additions of variances for multiphoton responses]

$$P(i|n=2) = \frac{1}{\sqrt{2\pi(\sigma_0^2 + 2\sigma_1^2)}} \exp\left[-\frac{(i - 2i_1)^2}{2(\sigma_0^2 + 2\sigma_1^2)}\right], \quad (15)$$

and more generally n photons,

$$P(i|n) = \frac{1}{\sqrt{2\pi(\sigma_0^2 + n\sigma_1^2)}} \exp\left[-\frac{(i - ni_1)^2}{2(\sigma_0^2 + n\sigma_1^2)}\right]. \quad (16)$$

Finally, since we know that the photon count n should be drawn out of the Poisson distribution, we can write

the expected distribution of currents as

$$P(i) = \sum_{n=0}^{\infty} P(i|n)P(n) \quad (17)$$

$$= \sum_{n=0}^{\infty} P(i|n)e^{-\bar{n}} \frac{\bar{n}^n}{n!} \quad (18)$$

$$= \sum_{n=0}^{\infty} \frac{\bar{n}^n}{n!} \frac{e^{-\bar{n}}}{\sqrt{2\pi(\sigma_0^2 + n\sigma_1^2)}} \exp\left[-\frac{(i - n\mu_1)^2}{2(\sigma_0^2 + n\sigma_1^2)}\right] \quad (19)$$

In Fig 4, we see that this really gives a very good description of the distribution that we observe when we sample the currents in response to a large number of flashes.

Problem 9: Exploring the sampling problem. The data that we see in Fig 4 are not a perfect fit to our model. On the other hand, there are only 350 samples that we are using to estimate the shape of the underlying probability distribution. This is an example of a problem that you will meet many times in comparing theory and experiment; perhaps you have some experience from physics lab courses which is relevant here. We will return to these issues of sampling and fitting nearer the end of the course, when we have some more powerful mathematical tools, but for now let me encourage you to play a bit. Use the model that leads to Eq (19) to generate samples of the peak current, and then use these samples to estimate the probability distribution. For simplicity, assume that $\mu_1 = 1$, $\sigma_0 = 0.1$, $\sigma_1 = 0.2$, and $\bar{n} = 1$. Notice that since the current is continuous, you have to make bins along the current axis; smaller bins reveal more structure, but also generate noisier results because the number of counts in each bin is smaller. As you experiment with different size bins and different numbers of samples, try to develop some feeling for whether the agreement between theory and experiment in Fig 4 really is convincing.

Seeing this distribution, and especially seeing analytically how it is constructed, it is tempting to draw lines along the current axis in the ‘troughs’ of the distribution, and say that (for example) when we observe a current of less than 0.5 pA, this reflects zero photons. Is this the right way for us—or for the toad’s brain—to interpret these data? To be precise, suppose that we want to set a threshold for deciding between $n = 0$ and $n = 1$ photon. Where should we put this threshold to be sure that we get the right answer as often as possible?

Suppose we set our threshold at some current $i = \theta$. If there really were zero photons absorbed, then if by chance $i > \theta$ we will incorrectly say that there was one photon. This error has a probability

$$P(\text{say } n = 1|n = 0) = \int_{\theta}^{\infty} di P(i|n = 0). \quad (20)$$

On the other hand, if there really was one photon, but by chance the current was less than the threshold, then we’ll say 0 when we should have said 1, and this has a probability

$$P(\text{say } n = 0|n = 1) = \int_{-\infty}^{\theta} di P(i|n = 1). \quad (21)$$

There could be errors in which we confuse two photons for zero photons, but looking at Fig 4 it seems that these higher order errors are negligible. So then the total probability of making a mistake in the $n = 0$ vs. $n = 1$ decision is

$$P_{\text{error}}(\theta) = P(\text{say } n = 1|n = 0)P(n = 0) + P(\text{say } n = 0|n = 1)P(n = 1) \quad (22)$$

$$= P(n = 0) \int_{\theta}^{\infty} di P(i|n = 0) + P(n = 1) \int_{-\infty}^{\theta} di P(i|n = 1). \quad (23)$$

We can minimize the probability of error in the usual way by taking the derivative and setting the result to zero at the optimal setting of the threshold, $\theta = \theta^*$:

$$\frac{dP_{\text{error}}(\theta)}{d\theta} = P(n = 0) \frac{d}{d\theta} \int_{\theta}^{\infty} di P(i|n = 0) + P(n = 1) \frac{d}{d\theta} \int_{-\infty}^{\theta} di P(i|n = 1) \quad (24)$$

$$= P(n = 0)(-1)P(i = \theta|n = 0) + P(n = 1)P(i = \theta|n = 1); \quad (25)$$

$$\left. \frac{dP_{\text{error}}(\theta)}{d\theta} \right|_{\theta=\theta^*} = 0 \Rightarrow P(i = \theta^*|n = 0)P(n = 0) = P(i = \theta^*|n = 1)P(n = 1). \quad (26)$$

This result has a simple interpretation. Given that we have observed some current i , we can calculate the probability that n photons were detected using Bayes’ rule for

conditional probabilities:

$$P(n|i) = \frac{P(i|n)P(n)}{P(i)}. \quad (27)$$

The combination $P(i|n=0)P(n=0)$ thus is proportional to the probability that the observed current i was generated by counting $n=0$ photons, and similarly the combination $P(i|n=1)P(n=1)$ is proportional to the probability that the observed current was generated by counting $n=1$ photons. The optimal setting of the threshold, from Eq (26), is when these two probabilities are equal. Another way to say this is that for each observable current i we should compute the probability $P(n|i)$, and then our “best guess” about the photon count n is the one which maximizes this probability. This guess is best in the sense that it minimizes the total probability of errors. This is how we draw the boundaries shown by dashed lines in Fig 4 [Check details! Also introduce names for these things—maximum likelihood, maximum a posteriori probability, This is also a place to anticipate the role of prior expectations in setting thresholds!]

Problem 10: More careful discrimination. You observe some variable x (e.g., the current flowing across the rod cell membrane) that is chosen either from the probability distribution $P(x|+)$ or from the distribution $P(x|-)$. Your task is to look at a particular x and decide whether it came from the $+$ or the $-$ distribution. Rather than just setting a threshold, as in the discussion above, suppose that when you see x you assign it to the $+$ distribution with a probability $p(x)$. You might think this is a good idea since, if you’re not completely sure of the right answer, you can hedge your bets by a little bit of random guessing. Express the




FIG. 5 Schematic of discrimination in the presence of noise. We have two possible signals, A and B, and we measure something, for example the current flowing across a cell membrane. Given either A or B, the current fluctuates. As explained in the text, the overall probability of confusing A with B is minimized if we draw a threshold at the point where the probability distributions cross, and identify all currents larger than this threshold as being B, all currents smaller than threshold as being A. Because the distributions overlap, it is not possible to avoid errors, and the area of the red shaded region counts the probability that we will misidentify A as B.

probability of a correct answer in terms of $p(x)$; this is a functional $P_{\text{correct}}[p(x)]$. Now solve the optimization problem for the function $p(x)$, maximizing P_{correct} . Show that the solution is deterministic [$p(x) = 1$ or $p(x) = 0$], so that if the goal is to be correct as often as possible you shouldn’t hesitate to make a crisp assignment even at values of x where you aren’t sure (!). Hint: Usually, you would try to maximize the P_{correct} by solving the variational equation $\delta P_{\text{correct}}/\delta p(x) = 0$. You should find that, in this case, this approach doesn’t work. What does this mean? Remember that $p(x)$ is a probability, and hence can’t take on arbitrary values.

Once we have found the decision rules that minimize the probability of error, we can ask about the error probability itself. As schematized in Fig 5, we can calculate this by integrating the relevant probability distributions on the ‘wrong sides’ of the threshold. For Fig 4, this error probability is less than three percent. Thus, under these conditions, we can look at the current flowing across the rod cell membrane and decide whether we saw $n=0, 1, 2 \dots$ photons with a precision such that we are wrong only on a few flashes out of one hundred. In fact, we might even be able to do better if instead of looking at the current at one moment in time we look at the whole trajectory of current vs. time, but to do this analysis we need a few more mathematical tools. Even without such a more sophisticated analysis, it’s clear that these cells really are acting as near perfect photon counters, at least over some range of conditions.

Problem 11: Asymptotic error probabilities. Should add a problem deriving the asymptotic probabilities of errors at high signal-to-noise ratios, including effects of prior probability.

A slight problem in our simple identification of the probability of seeing with the probability of counting K photons is that van der Velden found a threshold photon count of $K=2$, which is completely inconsistent with the $K=5-7$ found by Hecht, Shlaer and Pirenne. Barlow explained this discrepancy by noting that even when counting single photons we may have to discriminate (as in photomultipliers) against a background of dark noise.

Hecht, Shlaer and Pirenne inserted blanks in their experiments to be sure that you almost never say “I saw it” when nothing is there, which means you have to set a high threshold to discriminate against any background noise. On the other hand, van der Velden was willing to allow for some false positive responses, so his subjects could afford to set a lower threshold. Qualitatively, as shown in Fig 6, this makes sense, but to be a quantitative explanation the noise has to be at the right level.




FIG. 6 Trading of errors in the presence of noise. We observe some quantity that fluctuates even in the absence of a signal. When we add the signal these fluctuations continue, but the overall distribution of the observable is shifted. If set a threshold, declaring the signal is present whenever the threshold is exceeded, then we can trade between the two kinds of errors. At low thresholds, we never miss a signal, but there will be many false alarms. At high thresholds, there are few false alarms, but we miss most of the signals too. At some intermediate setting of the threshold, the total number of errors will be minimized.

One of the key ideas in the analysis of signals and noise is “referring noise to the input,” and we will meet this concept many times in what follows [more specific pointers]. Imagine that we have a system to measure something (here, the intensity of light, but it could be anything), and it has a very small amount of noise somewhere along the path from input to output. In many systems we will also find, along the path from input to output, an amplifier that makes all of the signals larger. But the amplifier doesn’t “know” which of its inputs are signal and which are noise, so everything is amplified. Thus, a small noise near the input can become a large noise near the output, but the size of this noise at the output does not, by itself, tell us how hard it will be to detect signals at the input. What we can do is to imagine that the whole system is noiseless, and that any noise we see at the output really was injected at the input, and thus followed exactly the same path as the signals we are trying to detect. Then we can ask how big this effective input noise needs to be in order to account for the output noise.

If the qualitative picture of Fig 6 is correct, then the minimum number of photons that we need in order to say “I saw it” should be reduced if we allow the observer the option of saying “I’m pretty sure I saw it,” in effect taking control over the trade between misses and false alarms. Barlow showed that this worked, quantitatively.

In the case of counting photons, we can think of the effective input noise as being nothing more than extra

“dark” photons, also drawn from a Poisson distribution. Thus if in the relevant window of time for detecting the light flash there are an average of 10 dark photons, for example, then because the variance of the Poisson distribution is equal to the mean, there will be fluctuations on the scale of $\sqrt{10}$ counts. To be very sure that we have seen something, we need an extra K real photons, with $K \gg \sqrt{10}$. Barlow’s argument was that we could understand the need for $K \sim 6$ in the Hecht, Shaler and Pirenne experiments if indeed there were a noise source in the visual system that was equivalent to counting an extra ten photons over the window in time and area of the retina that was being stimulated. What could this noise be?

In the frequency of seeing experiments, as noted above, the flash of light illuminated roughly 500 receptor cells on the retina, and subsequent experiments showed that one could find essentially the same threshold number of photons when the flash covered many thousands of cells. Furthermore, experiments with different durations for the flash show that human observers are integrating over ~ 0.1 s in order to make their decisions about whether they saw something. Thus, the “dark noise” in the system seems to be equivalent, roughly, to 0.1 photon per receptor cell per second, or less. To place this number in perspective, it is important to note that vision begins when the pigment molecule rhodopsin absorbs light and changes its structure to trigger some sequence of events in the receptor cell. We will learn much more about the dynamics of rhodopsin and the cascade of events responsible for converting this molecular event into electrical signals that can be transmitted to the brain, but for now we should note that if rhodopsin can change its structure by absorbing a photon, there must also be some (small) probability that this same structural change or “isomerization” will happen as the result of a thermal fluctuation. If this does happen, then it will trigger a response that is identical to that triggered by a real photon. Further, such rare, thermally activated events really are Poisson processes (see Section II.A), so that thermal activation of rhodopsin would contribute exactly a “dark light” of the sort we have been trying to estimate as a background noise in the visual system. But there are roughly one billion rhodopsin molecules per receptor cell, so that a dark noise of ~ 0.1 per second per cell corresponds to a rate of once per ~ 1000 years for the spontaneous isomerization of rhodopsin.

One of the key points here is that Barlow’s explanation works only if people actually can adjust the “threshold” K in response to different situations. The realization that this is possible was part of the more general recognition that detecting a sensory signal does not involve a true threshold between (for example) seeing and not seeing. Instead, *all* sensory tasks involve a discrimination between signal and noise, and hence there are different strategies which provide different ways of trading

off among the different kinds of errors. Notice that this picture matches what we know from the physics lab.

Problem 12: Simple analysis of dark noise. Suppose that we observe events drawn out of a Poisson distribution, and we can count these events perfectly. Assume that the mean number of events has two contributions, $\bar{n} = \bar{n}_{\text{dark}} + \bar{n}_{\text{flash}}$, where $\bar{n}_{\text{flash}} = 0$ if there is no light flash and $\bar{n}_{\text{flash}} = N$ if there is a flash. As an observer, you have the right to set a criterion, so that you declare the flash to be present only if you count $n \geq K$ events. As you change K , you change the errors that you make—when K is small you often say you saw something when nothing was there, but of hardly ever miss a real flash, while at large K the situation is reversed. The conventional way of describing this is to plot the fraction of “hits” (probability that you correctly identify a real flash) against the probability of a false alarm (i.e., the probability that you say a flash is present when it isn’t), with the criterion changing along the curve. Plot this “receiver operating characteristic” for the case $\bar{n}_{\text{dark}} = 10$ and $N = 10$. Hold \bar{n}_{dark} fixed and change N to see how the curve changes. Explain which slice through this set of curves was measured by Hecht et al, and the relationship of this analysis to what we saw in Fig 2.

There are classic experiments to show that people will adjust their thresholds automatically when we change the a priori probabilities of the signal being present, as expected for optimal performance. This can be done without any explicit instructions—you don’t have to tell someone that you are changing the probabilities—and it works in all sensory modalities, not just vision. At least implicitly, then, people learn something about probabilities and adjust their criteria appropriately. Threshold adjustments also can be driven by changing the rewards for correct answers or the penalties for wrong answers. In this view, it is likely that Hecht et al. drove their observers to high thresholds by having a large effective penalty for false positive detections. Although it’s not a huge literature, people have since manipulated these penalties and rewards in frequency of seeing experiments, with the expected results. Perhaps more dramatically, modern quantum optics techniques have been used to manipulate the statistics of photon arrivals at the retina, so that the tradeoffs among the different kinds of errors are changed ... again with the expected results.¹⁰

Not only did Baylor and coworkers detect the single photon responses from toad photoreceptor cells, they also found that single receptor cells in the dark show spontaneous photon-like events roughly at the right rate to be the source of dark noise identified by Barlow. If you




FIG. 7 [fill in the caption] From Aho et al (1988).

look closely you can find one of these spontaneous events in the earlier illustration of the rod cell responses to dim flashes, Fig 3. Just to be clear, Barlow identified a *maximum* dark noise level; anything higher and the observed reliable detection is impossible. The fact that the real rod cells have essentially this level of dark noise means that the visual system is operating near the limits of reliability set by thermal noise in the input. It would be nice to give a more direct test of this idea.

In the lab we often lower the noise level of photodetectors by cooling them. This should work in vision too, since one can verify that the rate of spontaneous photon-like events in the rod cell current is strongly temperature dependent, increasing by a factor of roughly four for every ten degree increase in temperature. Changing temperature isn’t so easy in humans, but it does work with cold blooded animals like frogs and toads. To set the stage, it is worth noting that one species of toad in particular (*Bufo bufo*) manages to catch its prey under conditions so dark that human observers cannot see the toad, much less the prey [add the reference!]. So, Aho et al. convinced toads to strike with their tongues at small worm-like objects illuminated by very dim lights, and measured how the threshold for reliable striking varied with temperature, as shown in Fig 7. Because one can actually make measurements on the retina itself, it is possible to calibrate light intensities as the rate at which rhodopsin molecules are absorbing photons and isomerizing, and the toad’s responses are almost deterministic once this rate is $r \sim 10^{-11} \text{ s}^{-1}$ in experiments at 15°C , and responses are detectable at intensities a factor of three to five below this. For comparison, the rate of thermal isomerizations at this temperature is $\sim 5 \times 10^{-12} \text{ s}^{-1}$.

If the dark noise consists of rhodopsin molecules spontaneously isomerizing at a rate r_d , then the mean number of dark events will be $n_d = r_d T N_r N_c$, where $T \sim 1 \text{ s}$ is the

¹⁰ It is perhaps too much to go through all of these results here, beautiful as they are. To explore, see the references at the end of the section.




FIG. 8 [fill in the caption] From Aho et al (1987, 1988).

relevant integration time for the decision, $N_r \sim 3 \times 10^9$ is the number of rhodopsin molecules per cell in this retina, and $N_c \sim 4,500$ is the number of receptor cells that are illuminated by the image of the worm-like object. Similarly, the mean number of real events is $n = rTN_rN_c$, and reliable detection requires $n > \sqrt{n_d}$, or

$$r > \sqrt{\frac{r_d}{TN_rN_c}} \sim 6 \times 10^{-13} \text{ s}^{-1}. \quad (28)$$

Thus, if the toad knows exactly which part of the retina it should be looking at, then it should reach a signal-to-noise ratio of one at light intensities a factor of ten below the nominal dark noise level. But there is no way to be sure where to look before the target appears, and the toad probably needs a rather higher signal-to-noise ratio before it is willing to strike. Thus it is plausible that the threshold light intensities in this experiment should be comparable to the dark noise level, as observed.

One can do an experiment very similar to the one with toads using human subjects (who say yes or no, rather than sticking out their tongues), asking for a response to small targets illuminated by steady, dim lights. Frogs will spontaneously jump at dimly illuminated patch of the ceiling, in an attempt to escape from an otherwise dark box. Combining these experiments, with the frogs held at temperatures from 10 to 20 °C, one can span a range of almost two orders of magnitude in the thermal isomerization rate of rhodopsin. It's not clear whether individual organisms hold their integration times T fixed as temperature is varied, or if the experiments on different organisms correspond to asking for integration over a similar total number of rhodopsin molecules (N_rN_c). Nonetheless, it is satisfying to see, in Fig 8, that the “threshold” light intensity, where response occur 50% of the time, is varying systematically with the dark noise level. It is certainly true that operating at lower temperatures allows

the detection of dimmer lights, or equivalently more reliable detection of the same light intensity,¹¹ as expected if the dominant noise source was thermal in origin. These experiments support the hypothesis that visual processing in dim lights really is limited by input noise and not by any inefficiencies of the brain.

Problem 13: Getting a feel for the brain's problem. Let's go back to Problem 3, where you simulated a Poisson process.

(a) If you use the strategy of making small bins $\Delta\tau$ and testing a random number in each bin against a threshold, then it should be no problem to generalize this to the case where the threshold is different at different times, so you are simulating a Poisson process in which the rate is varying as a function of time. As an example, consider a two second interval in which the counting rate has some background (like the dark noise in rods) value r_{dark} except in a 100 msec window where the rate is higher, say $r = r_{\text{dark}} + r_{\text{signal}}$. Remember that for one rod cell, r_{dark} is $\sim 0.02 \text{ sec}^{-1}$, while humans can see flashes which have $r_{\text{signal}} \sim 0.01 \text{ sec}^{-1}$ if they can integrate over 1000 rods. Try to simulate events in this parameter range and actually look at examples, perhaps plotted with x's to show you where the events occur on a single trial.

(b) Can you tell the difference between a trial where you have $r_{\text{signal}} = 0.01 \text{ sec}^{-1}$ and one in which $r_{\text{signal}} = 0$? Does it matter whether you know when to expect the extra events? In effect these plots give a picture of the problem that the brain has to solve in the Hecht-Shaler-Pirenne experiment, or at least an approximate picture.

(c) Sitting in a dark room to repeat the HSP experiment would take a long time, but maybe you can go from your simulations here to design a psychophysical experiment simple enough that you can do it on one another. Can you measure the reliability of discrimination between the different patterns of x's that correspond to the signal being present or absent? Do you see an effect of “knowing when to look”? Do people seem to get better with practice? Can you calculate the theoretical limit to how well one can do this task? Do people get anywhere near this limit? This is an open ended problem.

Problem 14: A better analysis? Go back to the original paper by Aho et (1988) and see if you can give a more compelling comparison between thresholds and spontaneous isomerization rates. From Eq (28), we expect that the light intensity required for some criterion level of reliability scales as the square root of the dark noise level, but also depends on the total number of rhodopsin molecules over which the subject must integrate. Can you estimate this quantity for the experiments on frogs and humans? Does this lead to an improved version of Fig 8? Again, this is an open ended problem.

The dominant role of spontaneous isomerization as a source of dark noise leads to a wonderfully counterintuitive result, namely that the photoreceptor which is

¹¹ The sign of this prediction is important. If we were looking for more reliable behaviors at higher temperatures, there could be many reasons for this, such as quicker responses of the muscles. Instead, the prediction is that we should see more reliable behavior as you cool down—all the way down to the temperature where behavior stops—and this is what is observed.

designed to maximize the signal-to-noise ratio for detection of dim lights will allow a significant number of photons to pass by undetected. Consider a rod photoreceptor cell of length ℓ , with concentration C of rhodopsin; let the absorption cross section of rhodopsin be σ . [Do I need to explain the definition of cross sections, and/or the derivation of Beer's law?] As a photon passes along the length of rod, the probability that it will be absorbed (and, presumably, counted) is $p = 1 - \exp(-C\sigma\ell)$, suggesting that we should make C or ℓ larger in order to capture more of the photons. But, as we increase C or ℓ , we are increasing the number of rhodopsin molecules, $N_{\text{rh}} = CA\ell$, with A the area of the cell, so we also increase the rate of dark noise events, which occurs at a rate r_{dark} per molecule.

If we integrate over a time τ , we will see a mean number of dark events (spontaneous isomerizations) $\bar{n}_{\text{dark}} = r_{\text{dark}}\tau N_{\text{rh}}$. The actual number will fluctuate, with a standard deviation $\delta n = \sqrt{\bar{n}_{\text{dark}}}$. On the other hand, if n_{flash} photons are incident on the cell, the mean number counted will be $\bar{n}_{\text{count}} = n_{\text{flash}}p$. Putting these factors together we can define a signal-to-noise ratio

$$SNR \equiv \frac{\bar{n}_{\text{count}}}{\delta n} = n_{\text{flash}} \frac{[1 - \exp(-C\sigma\ell)]}{\sqrt{CA\ell r_{\text{dark}}\tau}}. \quad (29)$$

The absorption cross section σ and the spontaneous isomerization rate r_{dark} are properties of the rhodopsin molecule, but as the rod cell assembles itself, it can adjust both its length ℓ and the concentration C of rhodopsin; in fact these enter together, as the product $C\ell$. When $C\ell$ is larger, photons are captured more efficiently and this leads to an increase in the numerator, but there also are more rhodopsin molecules and hence more dark noise, which leads to an increase in the denominator. Viewed as a function of $C\ell$, the signal-to-noise ratio has a maximum at which these competing effects balance; working out the numbers one finds that the maximum is reached when $C\ell \sim 1.26/\sigma$, and we note that all the other parameters have dropped out. In particular, this means that the probability of an incident photon *not* being absorbed is $1 - p = \exp(-C\sigma\ell) \sim e^{-1.26} \sim 0.28$. Thus, to maximize the signal-to-noise ratio for detecting dim flashes of light, nearly 30% of photons should pass through the rod without being absorbed (!). **Say something about how this compares with experiment!**

Problem 15: Escape from the tradeoff. Derive for yourself the numerical factor $(C\ell)_{\text{opt}} \sim 1.26/\sigma$. Can you see any way to design an eye which gets around this tradeoff between more efficient counting and extra dark noise? Hint: Think about what you see looking into a cat's eyes at night.




FIG. 9 Results of experiments in which observers are asked to rate the intensity of dim flashes, including blanks, on a scale from 0 to 6. Main figure shows that the variance of the ratings at fixed intensity is equal to the mean, as expected if the ratings are Poisson distributed. Insets show that the full distribution is approximately Poisson (upper) and that the mean rating is linearly related to the flash intensity, measured here as the mean number of photons delivered to the cornea. From Sakitt (1972).

If this is all correct, it should be possible to coax human subjects into giving responses that reflect the counting of individual photons, rather than just the summation of multiple counts up to some threshold of confidence or reliability. Suppose we ask observers not to say yes or no, but rather to rate the apparent intensity of the flash, say on a scale from 0 to 7. Remarkably, as shown in Fig 9, in response to very dim flashes interspersed with blanks, at least some observers will generate rating that, given the intensity, are approximately Poisson distributed: the variance of the ratings is essentially equal to the mean, and even the full distribution of ratings over hundreds of trials is close to Poisson. Further, the mean rating is linearly related to the light intensity, with an offset that agrees with other measurements of the dark noise level. Thus, the observers behaves exactly as if she can give a rating that is equal to the number of photons counted. This astonishing result would be almost too good to be true were it not that some observers deviate from this ideal behavior—they starting counting at two or three, but otherwise follow all the same rules.

While the phenomena of photon counting are very beautiful, one might worry that this represents just a very small corner of vision. Does the visual system continue to count photons reliably even when it's not completely dark outside? To answer this let's look at vision in a rather different animal, as in Fig 10. When you look down on the head of a fly, you see—almost to the exclusion of anything else—the large compound eyes. Each little hexagon that you see on the fly's head is a sepa-




FIG. 10 The fly’s eye(s). At left a photograph taken by H Leertouwer at the Rijksuniversiteit Groningen, showing (even in this poor reproduction) the hexagonal lattice of lenses in the compound eye. This is the blowfly *Calliphora vicina*. At right, a schematic of what a fly might see, due to Gary Larson. The schematic is incorrect because each lens actually looks in a different direction, so that whole eye (like ours) only has one image of the visual world. In our eye the “pixelation” of the image is enforced by the much less regular lattice of receptors on the retina; in the fly pixelation occurs already with the lenses.

rate lens, and in large flies there are $\sim 5,000$ lenses in each eye, with approximately 1 receptor cell behind each lens, and roughly 100 brain cells per lens devoted to the processing of visual information. The lens focuses light on the receptor, which is small enough to act as an optical waveguide. Each receptor sees only a small portion of the world, just as in our eyes; one difference between flies and us is that diffraction is much more significant for organisms with compound eyes—because the lenses are so small, flies have an angular resolution of about 1° , while we do about $100\times$ better. [Add figure to emphasize similarity of two eye types.]

The last paragraph was a little sloppy (“approximately one receptor cell”), so let’s try to be more precise. For flies there actually are eight receptors behind each lens. Two provide sensitivity to polarization and some color vision, which we will ignore here. The other six receptors look out through the same lens in different directions, but as one moves to neighboring lenses one finds that there is one cell under each of six neighboring lenses which looks in the same direction. Thus these six cells are equivalent to one cell with six times larger photon capture cross section, and the signals from these cells are collected and summed in the first processing stage (the lamina); one can even see the expected six fold improvement in signal to noise ratio, in experiments we’ll describe shortly.¹²

Because diffraction is such a serious limitation, one might expect that there would be fairly strong selection

¹² Talk about the developmental biology issues raised by these observations, and the role of the photoreceptors as a model system in developmental decision making. For example, Lubensky et al (2011). Not sure where to put this, though.

for eyes that make the most of the opportunities within these constraints. Indeed, there is a beautiful literature on optimization principles for the design of the compound eye; the topic even makes an appearance in Feynman’s undergraduate physics lectures. Roughly speaking (Fig 11), we can think of the fly’s head as being a sphere of radius R , and imagine that the lens are pixels of linear dimension d on the surface. Then the geometry determines an angular resolution (in radians) of $\delta\phi_{\text{geo}} \sim d/R$; resolution gets better if d gets smaller. On the other hand, diffraction through an aperture of size d creates a blur of angular width $\delta\phi_{\text{diff}} \sim \lambda/d$, where $\lambda \sim 500$ nm is the wavelength of the light we are trying to image; this limit of course improves as the aperture size d gets larger. Although one could try to give a more detailed theory, it seems clear that the optimum is reached when the two different limits are about equal, corresponding to an optimal pixel size

$$d_* \sim \sqrt{\lambda R}. \quad (30)$$

This is the calculation in the Feynman lectures, and Feynman notes that it gives the right answer within 10% in the case of a honey bee.

A decade before Feynman’s lectures, Barlow had derived the same formula and went into the drawers of the natural history museum in Cambridge to find a variety of insects with varying head sizes, and he verified that the pixel size really does scale with the square root of the head radius, as shown in Fig 12. I think this work should be more widely appreciated, and it has several features we might like to emulate. First, it explicitly brings measurements on many species together in a quantitative way. Second, the fact that multiple species can put




FIG. 11 At left, a schematic of the compound eye, with lenses of width d on the surface of a spherical eye with radius R . At right, the angular resolution of the eye as a function of the lens size, showing the geometric ($\delta\phi_{\text{geo}} \sim d/R$) and diffraction ($\delta\phi_{\text{diff}} \sim \lambda/d$) contributions in dashed lines; the full resolution in solid lines.

onto the same graph is not a phenomenological statement about, for example, scaling of one body part relative to another, but rather is based on a clearly stated physical principle. Finally, and most importantly for our later discussion in this course, Barlow makes an important transition: rather than just asking whether a biological system approaches the physical limits to performance, he assumes that the physical limits are reached and uses this hypothesis to predict something else about the structure of the system. This is, to be sure, a simple example, but an early and interesting example nonetheless.¹³




FIG. 12 The size of lenses in compound eyes as a function of head size, across many species of insect. From Barlow (1952).

[Should also point back to Mallock!]

Pushing toward diffraction-limited optics can't be the whole story, since at low light levels having lots of small pixels isn't much good—so few photons are captured in each pixel that one has a dramatic loss of intensity resolution. There must be some tradeoff between spatial resolution and intensity resolution, and the precise form of this tradeoff will depend on the statistical structure of the input images (if you are looking at clouds it will be different than looking at tree branches). The difficult question is how to quantify the relative worth of extra resolution in space vs intensity, and it has been suggested

¹³ This example also raises an interesting question. In Fig 12, each species of insect is represented by a single point. But not all members of the same species are the same size, as you must have noticed. Is the relationship between R and d that optimizes function preserved across the natural sizes variations among individuals? Does it matter whether the size differences are generated by environmental or genetic factors? This is a question about the reproducibility of spatial structures in development, a question we will come back to (albeit in simpler forms) in Section III.C. It would be good, though, if someone just *measured* the variations in eye dimensions across many individuals!

that the right way to do this is to count bits—design the eye not to maximize resolution, but rather to maximize the information that can be captured about the input image. This approach was a semi-quantitative success, showing how insects that fly late at night or with very high speeds (leading to blurring by photoreceptors with finite time resolution) should have less than diffraction limited spatial resolving power. I still think there are open questions here, however.

Coming back to the question of photon counting, one can record the voltage signals in the photoreceptor cells and detect single photon responses, as in vertebrates. If we want to see what happens at higher counting rates, we have to be sure that we have the receptor cells in a state where they don't "run down" too much because the increased activity. In particular, the rhodopsin molecule itself has to be recycled after it absorbs a photon. In animals with backbones, this actually happens not within the photoreceptor, but in conjunction with other cells that form the pigment epithelium. In contrast, in invertebrates the "resetting" of the rhodopsin molecule occurs within the receptor cell and can even be driven by absorption of additional long wavelength photons. Thus, if you want to do experiments at high photon flux on isolated vertebrate photoreceptors, there is a real problem of running out of functional rhodospin, but this doesn't happen in the fly's eye. Also, the geometry of the fly's eye makes it easier to do stable intracellular measurements without too much dissection.

To set the stage for experiments at higher counting rates, consider a simple model in which each photon arriving at time t_i produces a pulse $V_0(t - t_i)$, and these pulses just add to give the voltage [maybe there should be a sketch showing the summation of pulses to give the total voltage]

$$V(t) = V_{\text{DC}} + \sum_i V_0(t - t_i), \quad (31)$$

where V_{DC} is the constant voltage that one observes across the cell membrane in the absence of light. In Section A.1, we can find the distribution of the arrival times $\{t_i\}$ on the hypothesis that the photons arrive as a Poisson process with a time dependent rate $r(t)$; from Eq (A13) we have

$$P[\{t_i\}|r(t)] = \exp\left[-\int_0^T d\tau r(\tau)\right] \frac{1}{N!} r(t_1)r(t_2)\cdots r(t_N), \quad (32)$$

where $r(t)$ is the rate of photon arrivals—the light intensity in appropriate units. To compute the average voltage response to a given time dependent light intensity, we have to do a straightforward if tedious calculation:

$$\left\langle \sum_i V_0(t - t_i) \right\rangle = \sum_{N=0}^{\infty} \int_0^T d^N t_i P[\{t_i\}|r(t)] \sum_i V_0(t - t_i). \quad (33)$$

This looks a terrible mess. Actually, it's not so bad, and one can proceed systematically to do all of the integrals. Once you have had some practice, this isn't too difficult, but the first time through it is a bit painful, so I'll push the details off into Section A.1, along with all the other details about Poisson processes. When the dust settles [leading up to Eq (A64)], the voltage responds linearly to the light intensity,

$$\langle V(t) \rangle = V_{\text{DC}} + \int_{-\infty}^{\infty} dt' V_0(t-t') r(t'). \quad (34)$$

In particular, if we have some background photon counting rate \bar{r} that undergoes fractional modulations $C(t)$, so that

$$r(t) = \bar{r}[1 + C(t)], \quad (35)$$

then there is a linear response of the voltage to the contrast C ,

$$\langle \Delta V(t) \rangle = \bar{r} \int_{-\infty}^{\infty} dt' V_0(t-t') C(t'). \quad (36)$$

Recall that such integral relationships (convolutions) simplify when we use the Fourier transform. For a function of time $f(t)$ we will define the Fourier transform with the conventions

$$\tilde{f}(\omega) = \int_{-\infty}^{\infty} dt e^{+i\omega t} f(t), \quad (37)$$

$$f(t) = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} e^{-i\omega t} \tilde{f}(\omega). \quad (38)$$

Then, for two functions of time $f(t)$ and $g(t)$, we have

$$\int_{-\infty}^{\infty} dt e^{+i\omega t} \left[\int_{-\infty}^{\infty} dt' f(t-t') g(t') \right] = \tilde{f}(\omega) \tilde{g}(\omega). \quad (39)$$

Problem 16: Convolutions. Verify the “convolution theorem” in Eq (39). If you need some reminders, see, for example, Lighthill (1958).

Armed with Eq (39), we can write the response of the photoreceptor in the frequency domain,

$$\langle \Delta \tilde{V}(\omega) \rangle = \bar{r} \tilde{V}_0(\omega) \tilde{C}(\omega), \quad (40)$$

so that there is a transfer function, analogous to impedance relating current and voltage in an electrical circuit,

$$\tilde{T}(\omega) \equiv \frac{\langle \Delta \tilde{V}(\omega) \rangle}{\tilde{C}(\omega)} = \bar{r} \tilde{V}_0(\omega). \quad (41)$$

Recall that this transfer function is a complex number at every frequency, so it has an amplitude and a phase,

$$\tilde{T}(\omega) = |\tilde{T}(\omega)| e^{i\phi_T(\omega)}. \quad (42)$$

The units of \tilde{T} are simply voltage per contrast. The interpretation is that if we generate a time varying contrast $C(t) = C \cos(\omega t)$, then the voltage will also vary at frequency ω ,

$$\langle \Delta V(t) \rangle = |\tilde{T}(\omega)| C \cos[\omega t - \phi_T(\omega)]. \quad (43)$$

[Should we have one extra problem to verify this last equation? Or is it obvious?]

If every photon generates a voltage pulse $V_0(t)$, but the photons arrive at random, then the voltage must fluctuate. To characterize these fluctuations, we'll use some of the general apparatus of correlation functions and power spectra. A review of these ideas is given in Appendix A.2.

We want to analyze the fluctuations of the voltage around its mean, which we will call $\delta V(t)$. By definition, the mean of this fluctuation is zero, $\langle \delta V(t) \rangle = 0$. There is a nonzero variance, $\langle [\delta V(t)]^2 \rangle$, but to give a full description we need to describe the covariance between fluctuations at different times, $\langle \delta V(t) \delta V(t') \rangle$. Importantly, we are interested in systems that have no internal clock, so this covariance or correlation can't depend separately on t and t' , only on the difference. More formally, if *we* shift *our* clock by a time τ , this can't matter, so we must have

$$\langle \delta V(t) \delta V(t') \rangle = \langle \delta V(t + \tau) \delta V(t' + \tau) \rangle; \quad (44)$$

this is possible only if

$$\langle \delta V(t) \delta V(t') \rangle = C_V(t - t'), \quad (45)$$

where $C_V(t)$ is the “correlation function of V .” Thus, invariance under time translations restricts the form of the covariance. Another way of expressing time translation invariance in the description of random functions is to say that any particular wiggle in plotting the function is equally likely to occur at any time. This property also is called “stationarity,” and we say that fluctuations that have this property are stationary fluctuations.

In Fourier space, the consequence of invariance under time translations can be stated more simply—if we compute the covariance between two frequency components, we find

$$\langle \delta \tilde{V}(\omega_1) \delta \tilde{V}(\omega_2) \rangle = 2\pi \delta(\omega_1 + \omega_2) S_V(\omega_1), \quad (46)$$

where $S_V(\omega)$ is called the power spectrum (or power spectral density) of the voltage V . Remembering that $\delta \tilde{V}(\omega)$ is a complex number, it might be more natural to write this as

$$\langle \delta \tilde{V}(\omega_1) \delta \tilde{V}^*(\omega_2) \rangle = 2\pi \delta(\omega_1 - \omega_2) S_V(\omega_1). \quad (47)$$

Time translation invariance thus implies that fluctuations at different frequencies are independent.¹⁴ This makes sense, since if (for example) fluctuations at 2 Hz and 3 Hz were correlated, we could form beats between these components and generate a clock that ticks every second. Finally, the Wiener–Khinchine theorem states that the power spectrum and the correlation function are a Fourier transform pair,

$$S_V(\omega) = \int d\tau e^{+i\omega\tau} C_V(\tau), \quad (48)$$

$$C_V(\tau) = \int \frac{d\omega}{2\pi} e^{-i\omega\tau} S_V(\omega). \quad (49)$$

Notice that

$$\langle [\Delta V(t)]^2 \rangle \equiv C_V(0) = \int \frac{d\omega}{2\pi} S_V(\omega). \quad (50)$$

Thus we can think of each frequency component as having a variance $\sim S_V(\omega)$, and by summing these components we obtain the total variance.

Problem 17: More on stationarity. Consider some fluctuating variable $x(t)$ that depends on time, with $\langle x(t) \rangle = 0$. Show that, because of time translation invariance, higher order correlations among Fourier components are constrained:

$$\langle \tilde{x}(\omega_1) \tilde{x}^*(\omega_2) \tilde{x}^*(\omega_3) \rangle \propto 2\pi \delta(\omega_1 - \omega_2 - \omega_3) \quad (51)$$

$$\langle \tilde{x}(\omega_1) \tilde{x}(\omega_2) \tilde{x}^*(\omega_3) \tilde{x}^*(\omega_4) \rangle \propto 2\pi \delta(\omega_1 + \omega_2 - \omega_3 - \omega_4). \quad (52)$$

If you think of \tilde{x}^* (or \tilde{x}) as being analogous to the operators for creation (or annihilation) of particles, explain how these relations are related to conservation of energy for scattering in quantum systems.

Problem 18: Brownian motion in a harmonic potential. [The harmonic oscillator gets used more than once, of course; check for redundancy among problems in different sections!] Consider a particle of mass m hanging from a spring of stiffness κ , surrounded through a fluid. The effect of the fluid is, on average, to generate a drag force, and in addition there is a ‘Langevin force’ that describes the random collisions of the fluid molecules with the particle, resulting in Brownian motion. The equation of motion is

$$m \frac{d^2 x(t)}{dt^2} + \gamma \frac{dx(t)}{dt} + \kappa x(t) = \eta(t), \quad (53)$$

where γ is the drag coefficient and $\eta(t)$ is the Langevin force. A standard result of statistical mechanics is that the correlation function of the Langevin force is

$$\langle \eta(t) \eta(t') \rangle = 2\gamma k_B T \delta(t - t'), \quad (54)$$

where T is the absolute temperature and $k_B = 1.36 \times 10^{-23}$ J/K is Boltzmann’s constant.

(a.) Show that the power spectrum of the Langevin force is $S_\eta(\omega) = 2\gamma k_B T$, independent of frequency. Fluctuations with such a constant spectrum are called ‘white noise.’

(b.) Fourier transform Eq (53) and solve, showing how $\tilde{x}(\omega)$ is related to $\tilde{\eta}(\omega)$. Use this result to find an expression for the power spectrum of fluctuations in x , $S_x(\omega)$.

(c.) Integrate the power spectrum $S_x(\omega)$ to find the total variance in x . Verify that your result agrees with the equipartition theorem,

$$\left\langle \frac{1}{2} \kappa x^2 \right\rangle = \frac{1}{2} k_B T. \quad (55)$$

Hint: The integral over ω can be done by closing a contour in the complex plane.

(d.) Show that the power spectrum of the velocity, $S_v(\omega)$, is related to the power spectrum of position through

$$S_v(\omega) = \omega^2 S_x(\omega). \quad (56)$$

Using this result, verify the other prediction of the equipartition theorem for this system,

$$\left\langle \frac{1}{2} m v^2 \right\rangle = \frac{1}{2} k_B T. \quad (57)$$

Now we have a language for describing the signals and noise in the receptor cell voltage, by going to the frequency domain. What does this have to do with counting photons? The key point is that we can do a calculation similar to the derivation of Eq (40) for $\langle \Delta V(t) \rangle$ to show that, at $C = 0$, the voltage will undergo fluctuations—responding to the random arrival of photons—with power spectrum

$$N_V(\omega) = \bar{r} |V_0(\omega)|^2. \quad (58)$$

We call this N_V because it is noise. The noise has a spectrum shaped by the pulses V_0 , and the magnitude is determined by the photon counting rate; again see Appendix A.1 for details.

Notice that both the transfer function and noise spectrum depend on the details of $V_0(t)$. In particular, because this pulse has finite width in time, the transfer function gets smaller at higher frequencies. Thus if you watch a flickering light, the strength of the signal transmitted by your photoreceptor cells will decrease with increasing frequency.

The crucial point is that, for an ideal photon counter, although higher frequency signals are attenuated the signal-to-noise ratio actually doesn’t depend on frequency. Thus if we form the ratio

$$\frac{|\tilde{T}(\omega)|^2}{N_V(\omega)} = \frac{|\bar{r} \tilde{V}_0(\omega)|^2}{\bar{r} |V_0(\omega)|^2} = \bar{r}, \quad (59)$$

we just recover the photon counting rate, independent of details. Since this is proportional to the signal-to-noise ratio for detecting contrast modulations $\tilde{C}(\omega)$, we expect that real photodetectors will give less than this ideal value. [Should be able to make a crisper statement here—is it a theorem? Prove it, or give the proof as a problem?]

¹⁴ Caution: this is true only at second order; it is possible for different frequencies to be correlated when we evaluate products of three or more terms. See the next problem for an example.

Problem 19: Frequency vs counting rate. [Need to give more guidance through this problem! Step by step ...] If we are counting photons at an average rate \bar{r} , you might think that it is easier to detect variations in light intensity at a frequency $\omega \ll \bar{r}$ than at higher frequencies, $\omega \gg \bar{r}$; after all, in the high frequency case, the light changes from bright to dim and back even before (on average) a single photon has been counted. But Eq (59) states that the signal-to-noise ratio for detecting contrast in an ideal photon counter is independent of frequency, counter to this intuition. Can you produce a simple simulation to verify the predictions of Eq (59)? As a hint, you should think about observing the photon arrivals over a time T such that $\bar{r}T \gg 1$. Also, if you are looking for light intensity variations of the form $r(t) = \bar{r}[1 + C \cos(\omega t)]$, you should process the photon arrival times $\{t_i\}$ to form a signal $s = \sum_i \cos(\omega t_i)$.

So now we have a way of testing the photoreceptors: Measure the transfer function $\tilde{T}(\omega)$ and the noise spectrum $N_V(\omega)$, form the ratio $|\tilde{T}(\omega)|^2/N_V(\omega)$, and compare this with the actual photon counting rate \bar{r} . This was done for the fly photoreceptors, with the results shown in Fig 13. It's interesting to look back at the original papers and understand how they calibrated the measurement of \bar{r} (I'll leave this as an exercise for you!). [This account of the experiments is too glib. I will go back to expand and clarify. Rob has also offered new versions of the figures.]

What we see in Fig 13 is that, over some range of frequencies, the performance of the fly photoreceptors is close to the level expected for an ideal photon counter. It's interesting to see how this evolves as we change the mean light intensity, as shown in Fig 14. The performance of the receptors tracks the physical optimum up




FIG. 13 Signal and noise in fly photoreceptors, with experiments at four different mean light intensities, from de Ruyter van Steveninck & Laughlin (1996b). (a) Transfer function $|\tilde{T}(\omega)|^2$ from contrast to voltage. (b) Power spectrum of voltage noise, $N_V(\omega)$. (c) The ratio $|\tilde{T}(\omega)|^2/N_V(\omega)$, which would equal the photon counting rate if the system were ideal; dashed lines show the actual counting rates.




FIG. 14 Performance of fly photoreceptors vs light intensity. [Should redraw this, and label with consistent notation.] Having measured the quantity $\lambda_{\text{eff}} = |\tilde{T}(\omega)|^2/N_V(\omega)$, as in Fig 13, we plot the maximum value (typically at relatively low frequencies) vs the actual photon counting rate \bar{r} . We see that, over an enormous dynamic range, the signal-to-noise ratio tracks the value expected for an ideal photon counter.

to counting rates of $\bar{r} \sim 10^5$ photons/s. Since the integration time of the receptors is ~ 10 ms, this means that the cell can count, almost perfectly, up to about 1000.

An important point about these results is that they wouldn't work if the simple model were literally true. At low photon counting rates \bar{r} , the pulse V_0 has an amplitude of several millivolts, as you can work out from panel (a) in Fig 13. If we count $\sim 10^3$ events, this should produce a signal of several Volts, which is absolutely impossible in a real cell! What happens is that the system has an automatic gain control which reduces the size of the pulse V_0 as the light intensity is increased. Remarkably, this gain control or *adaptation* occurs while preserving (indeed, enabling) nearly ideal photon counting. Thus as the lights go up, the response to each photon become smaller (and, if you look closely, faster), but no less reliable.

Problem 20: Looking at the data. Explain how the data in Fig 13 provide evidence for the adaption of the pulse V_0 with changes in the mean light intensity.

[This seems a little brief! Maybe there should be a summary of what has happened, what we conclude ... also explain where the loose ends remain vs where things are solid.] These observations on the ability of the visual system to count single photons—down to the limit set by thermal noise in rhodopsin and up to counting rates of $\sim 10^5 \text{ s}^{-1}$ —raise questions at several different levels:

1. At the level of single molecules, we will see that the performance of the visual system depends crucially on the dynamics of rhodopsin itself. In particular, the structural response of the molecule to photon absorption is astonishingly fast, while the dark noise level means that the rate of spontaneous structural changes is extremely slow.

2. At the level of single cells, there are challenges in understanding how a network of biochemical reactions converts the structural changes of single rhodopsin molecules into macroscopic electrical currents across the rod cell membrane.

3. At the level of the retina as a whole, we would like to understand how these signals are integrated without being lost into the inevitable background of noise. Also at the level of the retina, we need to understand how single photon signals are encoded into the stereotyped pulses that are the universal language of the brain.

4. At the level of the whole organism, there are issues about how the brain learns to make the discriminations that are required for optimal performance.

In the next sections we'll look at these questions, in order.

It is a pleasure to read classic papers, and surely Hecht et al (1942) and van der Velden (1944) are classics, as is the discussion of dark noise by Barlow (1956). The pre-history of the subject, including the story about Lorentz, is covered by Bouman (1961). The general idea that our perceptual "thresholds" are really thresholds for discrimination against background noise with some criterion level of reliability made its way into quantitative psychophysical experiments in the 1950s and 60s, and this is now (happily) a standard part of experimental psychology; the canonical treatment is by Green and Swets (1966). The origins of these ideas are an interesting mix of physics and psychology, developed largely for radar in during World War II, and a summary of this early work is in the MIT Radar Lab series (Lawson & Uhlenbeck 1950). Another nice mix of physics and psychology is the revisiting of the original photon counting experiments using light sources with non-Poisson statistics (Teich et al 1982). The idea that random arrival of photons could limit our visual perception beyond the "just visible" was explored, early on, by de Vries (1943) and Rose (1948). Some of the early work by de Vries and coworkers on the physics of the sense organs (not just vision) is described in a lovely review (de Vries 1956). As a sociological note, de Vries was an experimental physicist with very broad interests, from biophysics to radiocarbon dating; for a short biography see de Waard (1960).

Barlow 1956: Retinal noise and absolute threshold. HB Barlow, *J Opt Soc Am* **46**, 634–639 (1956).

Bouman 1961: History and present status of quantum theory in vision. MA Bouman, in *Sensory Communication*, W Rosenblith, ed, pp 377–401 (MIT Press, Cambridge, 1961).

Green & Swets 1966: *Signal Detection Theory and Psychophysics*. DM Green, & JA Swets (Wiley, New York, 1966).

Hecht et al 1942: Energy, quanta and vision. S Hecht, S Schlaer & MH Pirenne, *J Gen Physiol* **25**, 819–840 (1942).

Lawson & Uhlenbeck 1950: *Threshold Signals*. MIT Radiation Laboratory Series vol 24 JL Lawson & GE Uhlenbeck (McGraw-Hill, New York, 1950).

Rose 1948: The sensitivity performance of the human eye on an absolute scale. A Rose, *J Opt Soc Am* **38**, 196–208 (1948).

Teich et al 1982: Multiplication noise in the human visual system at threshold. III: The role of non-Poisson quantum fluctuations. MC Teich, PR Prucnal, G Vannucci, ME Breton & WJ McGill, *Biol Cybern* **44**, 157–165 (1982).

van der Velden 1944: Over het aantal lichtquanta dat nodig is voor een lichtprikkel bij het menselijk oog. HA van der Velden, *Physica* **11**, 179–189 (1944).

de Vries 1943: The quantum character of light and its bearing upon threshold of vision, the differential sensitivity and visual acuity of the eye. Hl de Vries, *Physica* **10**, 553–564 (1943).

de Vries 1956: Physical aspects of the sense organs. Hl de Vries, *Prog Biophys Biophys Chem* **6**, 207–264 (1956).

de Waard 1960: Hessel de Vries, physicist and biophysicist. H de Waard, *Science* **131**, 1720–1721 (1960).

Single photon responses in receptor cells of the horseshoe crab were reported by Fuortes and Yeandle (1964). The series of papers from Baylor and co-workers on single photon responses in vertebrate rod cells, first from toads and then from monkeys, again are classics, well worth reading today, not least as examples of how to do quantitative experiments on biological systems. Aho, Donner, Reuter and co-workers have made a major effort to connect measurements on rod cells and ganglion cells with the behavior of the whole organism, using the toad as an example; among their results are the temperature dependence of dark noise (Fig 8), and the latency/anticipation results in Section I.D. The remarkable experiments showing that people really can count every photon are by Sakitt (1972). We will learn more about currents and voltages in cells very soon, but for background I have always liked Aidley's text, now in multiple editions; as is often the case, the earlier editions can be clearer and more compact.

Aidley 1998: *The Physiology of Excitable Cells, 4th Edition* DJ Aidley (Cambridge University Press, Cambridge, 1998).

Aho et al 1987: Retinal noise, the performance of retinal ganglion cells, and visual sensitivity in the dark-adapted frog. A-C Aho, K Donner, C Hydén, T Reuter & OY Orlov, *J Opt Soc Am A* **4**, 2321–2329 (1987).

Aho et al 1988: Low retinal noise in animals with low body temperature allows high visual sensitivity. A-C Aho, K Donner, C Hydén, LO Larsen & T Reuter, *Nature* **334**, 348–350 (1988).

Aho et al 1993: Visual performance of the toad (*Bufo bufo*) at low light levels: retinal ganglion cell responses and prey-catching accuracy. A-C Aho, K Donner, S Helenius, LO Larsen & T Reuter, *J Comp Physiol A* **172**, 671–682 (1993).

Baylor et al 1979a: The membrane current of single rod outer segments. DA Baylor, TD Lamb & K-W Yau, *J Physiol (Lond)* **288**, 589–611 (1979).

Baylor et al 1979b: Rod responses to single photons. DA Baylor, TD Lamb & K-W Yau, *J Physiol (Lond)* **288**, 613–634 (1979).

Baylor et al 1980: Two components of electrical dark noise in toad retinal rod outer segments. DA Baylor, G Matthews & K-W Yau, *J Physiol (Lond)* **309**, 591–621 (1980).

Baylor et al 1984: The photocurrent, noise and spectral sensitivity of rods of the monkey *Macaca fascicularis*. DA Baylor, BJ Nunn & JF Schnapf, *J Physiol (Lond)* **357**, 575–607 (1984).

Fuortes & Yeandle 1964: Probability of occurrence of discrete potential waves in the eye of *Limulus*. MGF Fuortes & S Yeandle, *J Gen Physiol* **47**, 443–463 (1964).

Sakitt 1972: Counting every quantum. B Sakitt, *J Physiol* **223**, 131–150 (1972).

For the discussion of compound eyes, useful background is contained in Stavenga and Hardie (1989), and in the beautiful compilation of insect brain anatomy by Strausfeld (1976), although this is hard to find; as an alternative there is an online atlas, <http://flybrain.neurobio.arizona.edu/Flybrain/html/>. There is also the more recent Land & Nilsson (2002). Evidently Larson (2003) is an imperfect guide to these matters. Everyone should have a copy of the Feynman lectures (Feynman et al 1963), and check the chapters on vision. The early work by Barlow (1952) deserves more appreciation, as noted in the main text, and the realization that diffraction must be important for insect eyes goes back to Mallock (1894). For a gentle introduction to the wider set of ideas about scaling relations between different body parts, see McMahon & Bonner (1983). The experiments on signal-to-noise ratio in fly photoreceptors are by de Ruyter van Steveninck and Laughlin (1996a, 1996b). For a review of relevant ideas in Fourier analysis and related matters, see Appendix A.2 and Lighthill (1958). You should come back to the ideas of Snyder et al (Snyder 1977, Snyder et al 1977) near the end of the book, after we have covered some of the basics of information theory.

Barlow 1952: The size of ommatidia in apposition eyes. HB Barlow, *J Exp Biol* **29**, 667–674 (1952).

Feynman et al 1963: *The Feynman Lectures on Physics*. RP Feynman, RB Leighton & M Sands (Addison-Wesley, Reading, 1963).

Larson 2003: *The Complete Far Side*. G Larson (Andrews McNeel Publishing, Kansas City, 2003).

Land & Nilsson 2002: *Animal Eyes* MF Land & D-E Nilsson (Oxford University Press, Oxford, 2002).

Lighthill 1958: *Introduction to Fourier Analysis and Generalized Functions*. MJ Lighthill (Cambridge University Press, Cambridge, 1958)

Mallock 1894: Insect sight and the defining power of composite eyes. A Mallock, *Proc R Soc Lond* **55**, 85–90 (1894).

McMahon & Bonner 1983: *On Size and Life*. TA McMahon & JT Bonner (WH Freeman, New York, 1983).

de Ruyter van Steveninck & Laughlin 1996a: The rate of information transfer at graded-potential synapses. RR de Ruyter van Steveninck & SB Laughlin, *Nature* **379**, 642–645 (1996).

de Ruyter van Steveninck & Laughlin 1996b: Light adaptation and reliability in blowfly photoreceptors. R de Ruyter van Steveninck & SB Laughlin, *Int. J. Neural Syst.* **7**, 437–444 (1996)

Snyder 1977: Acuity of compound eyes: Physical limitations and design. AW Snyder, *J Comp Physiol* **116**, 161–182 (1977).

Snyder et al 1977: Information capacity of compound eyes. AW Snyder, DS Stavenga & SB Laughlin, *J Comp Physiol* **116**, 183–207 (1977).

Stavenga & Hardie 1989: *Facets of Vision*. DG Stavenga & RC Hardie, eds (Springer-Verlag, Berlin, 1989).

Strausfeld 1976: *Atlas of an Insect Brain*. N Strausfeld (Springer-Verlag, Berlin, 1976).

Finally, a few reviews that place the results on photon counting into a broader context.

Barlow 1981: Critical limiting factors in the design of the eye and visual cortex. HB Barlow, *Proc R Soc Lond Ser B* **212**, 1–34 (1981).

Bialek 1987: Physical limits to sensation and perception. W Bialek, *Ann Rev Biophys Biophys Chem* **16**, 455–478 (1987).

Rieke & Baylor 1998: Single photon detection by rod cells of the retina. F Rieke & DA Baylor, *Rev Mod Phys* **70**, 1027–1036 (1998).




FIG. 15 Schematic structure of rhodopsin, showing the organic pigment retinal nestled in a pocket formed by the surrounding opsin protein. This conformation of the retinal is called 11-cis, since there is a rotation around the bond between carbons numbered 11 and 12 (starting at the lower right in the ring). Insets illustrate the conventions in such chemical structures, with carbons at nodes of the skeleton, and hydrogens not shown, but sufficient to make sure that each carbon forms four bonds.

B. Single molecule dynamics

To a remarkable extent, our ability to see in the dark is limited by the properties of rhodopsin itself, essentially because everything else works so well. Rhodopsin consists of a medium sized organic pigment, retinal, enveloped by a large protein, opsin (cf Fig 15). The primary photo-induced reaction is isomerization of the retinal, which ultimately couples to structural changes in the protein. The effort to understand the dynamics of these processes goes back to Wald's isolation of retinal (a vitamin A derivative) in the 1930s, his discovery of the isomerization, and the identification of numerous states through which the molecule cycles. The field was given a big boost by the discovery that there are bacterial rhodopsins, some of which serve a sensory function while others are energy transducing molecules, using the energy of the absorbed photon to pump protons across the cell membrane; the resulting difference in electrochemical potential for protons is a universal intermediate in cellular energy conversion, not just in bacteria but in us as well. [Maybe a pointer to channel rhodopsins would be good here too.]

By now we know much more than Wald did about the structure of the rhodopsin molecule [need to point to a better figure, more details].

While there are many remarkable features of the rhodopsin molecule, we would like to understand those particular features that contribute to the reliability of photon counting. First among these is the very low spontaneous isomerization rate, roughly once per thousand years. As we have seen, these photon-like events provide




FIG. 16 Isomerization of retinal, the primary event at the start of vision. The π -bonds among the carbons favor planar structures, but there are still alternative conformations. The 11-cis conformation is the ground state of rhodopsin, and after photon absorption the molecule converts to the all-trans configuration. These different structures have different absorption spectra, as well as other, more subtle differences. Thus we can monitor the progress of the transition 11-cis \rightarrow all-trans essentially by watching the molecule change color, albeit only slightly. [Show the spectra!]

the dominant noise source that limits our ability to see in the dark, so there is a clear advantage to having the lowest possible rate. When we look at the molecules themselves, purified from the retina, we can “see” the isomerization reaction because the initial 11-cis state and the final all-trans states (see Fig 16) have different absorption spectra [add this to the figure]. For rhodopsin itself, the spontaneous isomerization rate is too slow to observe in a bulk experiment. If we isolate the pigment retinal, however, we find that it has a spontaneous isomerization rate of $\sim 1/\text{yr}$, so that a bottle of 11-cis retinal is quite stable, but the decay to all-trans is observable.

How can we understand that rhodopsin has a spontaneous isomerization rate 1000 \times less than that of retinal? The spontaneous isomerization is thermally activated, and has a large “activation energy” as estimated from the temperature dependence of the dark noise.¹⁵ It seems reasonable that placing the retinal molecule into the pocket formed by the protein opsin would raise the activation energy, essentially because parts of the protein need to be pushed out of the way in order for the retinal to rotate and isomerize. Although this sounds plausible, it’s probably wrong. If we write the

¹⁵ I am assuming here that the ideas of activation energy and Arrhenius behavior of chemical reaction rates are familiar. For more on this, see Section II.A.

dark isomerization rate as $r = Ae^{-E_{\text{act}}/k_B T}$, retinal and rhodopsin have the same value of the activation energy $E_{\text{act}} = 21.9 \pm 1.6 \text{ kcal/mole}$ [this is from measurements on rods; give the number in useful units! maybe footnote about difficulties of units] within experimental error, but different values of the prefactor A . If we look at photoreceptor cells that are used for daytime vision—the cones, which also provide us with sensitivity to colors, as discussed below [check where this gets done!]¹⁶—the dark noise level is higher (presumably single photon counting is unnecessary in bright light), but again this is a difference in the prefactor, not in the activation energy. As we will see when we discuss the theory of reaction rates in Section II.A, understanding prefactors is much harder than understanding activation energies, and I think we don’t really have a compelling theoretical picture that explains the difference between retinal and rhodopsin. [Fred Rieke gave me some pointers I have to chase down before deciding on that last sentence!]

The isolated retinal pigment isomerization at a rate that is faster than rhodopsin. On the other hand, if we excite the isolated retinal with a very short pulse of light, and follow the resulting changes in absorption spectrum, these photo-induced dynamics are not especially fast, with isomerization occurring at a rate $\sim 10^9 \text{ s}^{-1}$. Although this is fast compared to the reactions that we can see directly, it is actually so slow that it is comparable to the rate at which the molecule will re-emit the photon. We recall from quantum mechanics that the spontaneous emission rates from electronic excited states are constrained by sum rules if they are dipole-allowed. This means that emission lifetimes for visible photons are order 1 nanosecond for almost all of the simple cases. In a big molecule, there can be some re-arrangement of the molecular structure before the photon is emitted (see the discussion below), and this results in the emitted or fluorescent photon being of longer wavelength. Nonetheless, the natural time scale is nanoseconds, and the isomerization of retinal is not fast enough to prevent fluorescence and truly capture the energy of the photon with high probability.

Problem 21: Why nanoseconds? Explain why spontaneous emission of visible photons typically occurs with a rate $\sim 10^9 \text{ s}^{-1}$. [Need to explain where to start!]

Now fluorescence is a disaster for visual pigment—not only don’t you get to count the photon where it was absorbed, it might get counted somewhere else, blurring the image. In fact rhodopsin does not fluoresce: The quantum yield or branching ratio for fluorescence is $\sim 10^{-5}$.




FIG. 17 [This needs to be redrawn; maybe two figures to make different points? Convert all the units once and for all?] Femtosecond dynamics of rhodopsin, from Wang et al (1994). At left, schematic potential energy surfaces in the electronic ground and excited states. At right, panel (A) shows transient absorption spectra following a 35 fs pulse of 500 nm light. Panel (B) shows the magnitude of the Fourier transform of the time dependent absorption at each of several wavelengths, illustrating the oscillations expected if the vibrational dynamics is coherent. You might like to convert the kcal/mol and cm^{-1} into more conventional physical units!

If we imagine the molecule sitting in the excited state, transitioning to the ground state via fluorescence at a rate $\sim 10^9 \text{ s}^{-1}$, then to have a branching ratio of 10^{-5} the competing process must have a rate of $\sim 10^{14} \text{ s}^{-1}$. Thus, the rhodopsin molecule must leave the excited state by some process on a time scale of ~ 10 femtoseconds, which is extraordinarily fast. Indeed, for many years, every time people built faster pulsed lasers, they went back to rhodopsin to look at the initial events, culminating in the direct demonstration of femtosecond isomerization, making this one of the fastest molecular events ever observed.

The 11-cis and all trans configurations of retinal have different absorption spectra, and this is why we can observe the events following photon absorption as an evolution of the spectrum. The basic design of such experiments is to excite the molecules with a brief pulse of light, elevating them into the excited state, and then probe with another brief pulse after some delay. In the simplest version, one repeats the experiment many times with different choices of the delay and the energy or wavelength of the probe pulse. An example of the results from such an experiment are shown in Fig 17. The first thing to notice is that the absorption at a wavelength of 550 nm, characteristic of the all-trans structure, rises very quickly after the pulse which excites the system, certainly within tens of femtoseconds. In fact this experiment reveals all sorts of interesting structure, to which we will return below.

The combination of faster photon induced isomerization and slower thermal isomerization means that the

protein opsin acts as an electronic state selective catalyst: ground state reactions are inhibited, excited state reactions accelerated, each by orders of magnitude. It is fair to say that if these state dependent changes in reaction rate did not occur—that is, if the properties of rhodopsin were those of retinal—then we simply would not be able to see in the dark of night.

Problem 22: What would vision be like if ... ? Imagine that the spontaneous isomerization rate and quantum yield for photoisomerization in rhodopsin were equal to those in retinal. Estimate, quantitatively, what this would mean for our ability to see at night. [we should try to connect with real intensities at dusk etc.]

In order to make sense out of all of this, and get started in understanding how rhodopsin achieves its function, we need to understand something about electronic transitions in large molecules, as opposed to the case of atoms that we all learned about in our quantum mechanics classes. The absorption of a photon by an atom involves a transition between two electronic states, and this is also true for a large molecule. But for the atom the absorption line is very narrow, while for big molecules it is very broad. For rhodopsin, there is a nice way of measuring the absorption spectrum over a very large dynamic range, and this is to use the rod cell as a sensor. Instead of asking how much light is absorbed, we can try assuming¹⁶ that all absorbed photons have a constant probability of generating a pulse of current at the rod's output, and so we can adjust the light intensity at each wavelength to produce the same current. If the absorption is stronger, we need less light, and conversely more light if the absorption is weaker. The results of such an experiment are shown in Fig 18. It is beautiful that in this way one can follow the long wavelength tail of the spectrum down to cross-sections that are $\sim 10^{-5}$ of the peak. More qualitatively, we see that the width of the spectrum, say at half maximum, is roughly 20% of the peak photon energy, which is enormous in contrast with atomic absorption lines.

As an aside, the fact that one can follow the sensitivity of the photoreceptor cell deep into the long wavelength tail opens the possibility of asking a very different question about the function of these cells (and all cells). We recall that every cell in our bodies has the same genetic

¹⁶ This assumption can also be checked. It's true, but I think there have not been very careful measurements in the long wavelength tail, where something interesting might happen.

material, and hence the instructions for making all possible proteins. In particular, all photoreceptor cells have the ability to make all visual pigments. But the different classes of receptors—rods and the three kinds of cones—make different pigments, corresponding to different proteins surrounding more or less the same retinal molecule, and the resulting differences in absorption spectra provide the basis for color vision. If a single cone couldn't reliably turn on the expression of one rhodopsin gene, and turn off all of the others, then the retina wouldn't be able to generate a mix of spectral sensitivities, and we wouldn't see colors. But how “off” is “off”?

In a macaque monkey (not so different from us in these matters), “red” cones have their peak sensitivity at a wavelength ~ 570 nm, but at this wavelength the “blue” cones have sensitivities that are $\sim 10^5 \times$ reduced relative to their own peak. Since the peak absorption cross-sections are comparable, this tells us that the relative concentration of red pigments in the blue cones must be less than 10^{-5} . That is, the cell makes at least 10^5 times as much of the correct protein as it does of the incorrect proteins, which I always thought was pretty impressive.¹⁷

Returning to the absorption spectrum itself, we realize that a full treatment would describe molecules by doing




FIG. 18 Sensitivity of the rod photoreceptor as a function of wavelength. This is measured, as explained in the text, by adjusting the intensity of light to give a criterion output, so that very low sensitivity corresponds to shining a bright light, rather than measuring a small output. Redrawn from Baylor et al (1979a).

¹⁷ Many thanks to Denis Baylor for reminding me of this argument. Since there are $\sim 10^9$ rhodopsins in one cell, errors of even one part in 10^5 would mean that there are thousands of “wrong” molecules floating around. I wonder if this is true, or if the true errors are even smaller. [Apparently there is evidence that some cones are less precise about what defines “off;” should check this!]




FIG. 19 Schematic of the electronic states in a large molecule, highlighting their coupling to motion of the nuclei. The sketch shows two states, with photon absorption (in blue) driving transitions between them. If we think in semi-classical terms, as explained in the text, then these transitions are ‘too fast’ for the atoms to move, and hence are vertical on such plots (the Franck–Condon approximation). Because the atomic coordinates fluctuate, as indicated by the Boltzmann distribution, the energy of the photon required to drive the transition also fluctuates, and this broadens the absorption spectrum.

the quantum mechanics of a combined system of electrons and nuclei. But the nuclei are very much heavier than the electrons, and hence move more slowly. More rigorously, the large ratio of masses means that we can think of solving the quantum mechanics of the electrons with the nuclei in fixed position, and then for each such atomic configuration the energy of the electrons contributes to the potential energy; as the nuclei move in this potential (whether classically or quantum mechanically) the electrons follow adiabatically.¹⁸ This is the Born–Oppenheimer approximation, which is at the heart of all attempts to understand molecular dynamics.¹⁹

Figure 19 shows the energy of two different electronic states, plotted schematically against (one of the) atomic coordinates. In the ground state, we know that there is some arrangement of the atoms that minimizes the en-

¹⁸ Because the electrons (mostly) follow the nuclei, I will use “nuclei” and “atoms” interchangeably in what follows.

¹⁹ I assume that most readers know something about the Born–Oppenheimer approximation, since it is a pretty classical subject. It is also one of the first adiabatic approximations in quantum mechanics. It took many years to realize that some very interesting things can happen in the adiabatic limit, notably the appearance of non-trivial phase factors in the adiabatic evolution of wave functions. Some of these ‘complications’ (to use a word from one of original papers) were actually discovered in the context of the Born–Oppenheimer approximation itself, but now we know that this circle of ideas is much bigger, extending out to quantum optics and quite exotic field theories.

ergy, and that in the neighborhood of this minimum the potential surface must look roughly like that of a system of Hookean springs. Once we lift the electrons into the first excited state, there is again some configuration of the atoms that minimizes the energy (unless absorbing one photon is enough to break the molecule apart!), but unless there is some symmetry this equilibrium configuration will be different than in the ground state, and the stiffness of the spring holding the molecule in this equilibrium configuration also will be different. Hence in Fig 19, the energy surfaces for the ground and excited states are shown displaced and with different curvatures.

It is important to realize that sketches such as that in Fig 19 are approximations in many senses. Most importantly, this sketch involves only one coordinate. You may be familiar with a similar idea in the context of chemical reactions, where out of all the atoms that move during the reaction we focus on one “reaction coordinate” that forms a path from the reactants to products; for more about this see Section II.A. One view is that this is just a convenience—we can’t draw in many dimensions, so we just draw one, and interpret the figure cautiously. Another view is that the dynamics *are* effectively one dimensional, either because there is a separation of time scales, or because we can change coordinates to isolate, for example, a single coordinate that couples to the difference in energy between the ground and excited electronic states. The cost of this reduction in dimensionality might be a more complex dynamics along this one dimension, for example with a “viscosity” that is strongly frequency dependent, which again means that we need to be cautious in interpreting the picture that we draw. In what follows I’ll start by being relatively informal, and try to become more precise as we go along.

In the limit that the atoms are infinitely heavy, they don’t move appreciably during the time required for an electronic transition. On the other hand, the positions of the atoms still have to come out of the Boltzmann distribution, since the molecule is in equilibrium with its environment at temperature T . In this limit, we can think of transitions between electronic states as occurring without atomic motion, corresponding to vertical lines on the schematic in Fig 19. If the photon happens to arrive when the atomic configuration is a bit to the left of the equilibrium point, then as drawn the photon energy needs to be larger in order to drive the transition; if the configuration is a bit to the right, then the photon en-

ergy is smaller. In this way, the Boltzmann distribution of atomic positions is translated into a broadening of the absorption line. In particular, the transition can occur with a photon that has very little energy if we happen to catch a molecule in the rightward tail of the Boltzmann distribution: the electronic transition can be made up partly from the energy of the photon and partly from energy that is “borrowed” from the thermal bath. As a result, the absorption spectrum should have a tail at long wavelengths, and this tail will be strongly temperature dependent, and this is observed in rhodopsin and other large molecules. Since our perception of color depends on the relative absorption of light by rhodopsins with different spectra, this means that there must be wavelengths such that the apparent color of the light will depend on temperature [need a pointer and refs for this .. maybe tell the story of de Vries and the hot tub?]

Concretely, if we imagine that the potential surfaces are perfect Hookean springs, but with displaced equilibrium positions, then we can relate the width of the spectrum directly to the magnitude of this displacement. In the ground state we have the potential

$$V_g(q) = \frac{1}{2}\kappa q^2, \quad (60)$$

and in the excited state we have

$$V_e(q) = \epsilon + \frac{1}{2}\kappa(q - \Delta)^2, \quad (61)$$

where ϵ is the minimum energy difference between the two electronic states and Δ is the shift in the equilibrium position, as indicated in Fig 20. With q fixed, the condition for absorbing a photon is that the energy $\hbar\Omega$ match the difference in electronic energies,

$$\hbar\Omega = V_e(q) - V_g(q) = \epsilon + \frac{1}{2}\kappa\Delta^2 - \kappa\Delta q. \quad (62)$$

The probability distribution of q when molecules are in the ground state is given by

$$P(q) = \frac{1}{Z} \exp\left[-\frac{V_g(q)}{k_B T}\right] = \frac{1}{\sqrt{2\pi\kappa k_B T/\kappa}} \exp\left[-\frac{\kappa q^2}{2k_B T}\right], \quad (63)$$

so we expect the cross-section for absorbing a photon of frequency Ω to have the form

$$\sigma(\Omega) \propto \int dq P(q) \delta\left[\hbar\Omega - \left(\epsilon + \frac{1}{2}\kappa\Delta^2 - \kappa\Delta q\right)\right] \quad (64)$$

$$\propto \int dq \exp\left[-\frac{\kappa q^2}{2k_B T}\right] \delta\left[\hbar\Omega - \left(\epsilon + \frac{1}{2}\kappa\Delta^2 - \kappa\Delta q\right)\right] \quad (65)$$

$$\propto \exp\left[-\frac{(\hbar\Omega - \hbar\Omega_{\text{peak}})^2}{4\lambda k_B T}\right], \quad (66)$$

where the peak of the absorption is at

$$\hbar\Omega_{\text{peak}} = \epsilon + \lambda, \quad (67)$$

and

$$\lambda = \frac{1}{2}\kappa\Delta^2 \quad (68)$$

is the energy required to distort the molecule into the equilibrium configuration of the excited state if we stay in the ground state.




FIG. 20 The potential surfaces of Fig 19, redrawn in the special case where they are parabolic. Then, as in Eqs (60) through (68), there are just a few key parameters that determine the shape of the absorption spectrum and also the fluorescence emission. **Redraw figure to show that $\hbar\Omega_{\text{peak}} = \epsilon + \lambda$; ref to Eq (67).**

The energy λ is known, in different contexts, as the reorganization energy or the Stokes shift. If the molecule stays in the excited state for a long time, the distribution of coordinates will re-equilibrate to the Boltzmann distribution appropriate to $V_e(q)$, so that the most likely coordinate becomes $q = \Delta$. At this coordinate, if the molecule returns to the ground state by emitting a photon—fluorescence—the energy of this photon will be $\hbar\Omega_{\text{fluor}} = \epsilon - \lambda$. Thus the peak fluorescence is at lower energies, or red shifted from the absorption peak by an amount 2λ , as one can read off from Fig 20. This connects the width of the absorption band to the red shift that occurs in fluorescence, and for many molecules this prediction is correct, quantitatively, giving us confidence in the basic picture. **[I wonder if all of this needs more figures in order to be clear?]**

In the case of rhodopsin, the peak absorption is at a wavelength of 500 nm or an energy of $\hbar\Omega_{\text{peak}} = 2.5$ eV. The width of the spectrum is described roughly by a Gaussian with a standard deviation of $\sim 10\%$ of the peak energy, so that $2\lambda k_B T \sim (0.25 \text{ eV})^2$, or $\lambda \sim 1.25$ eV. Surely we can't take this seriously, since this reorganization energy is enormous, and would distort the molecule well beyond the point where we could describe the potential surfaces by Hookean springs. Amusingly, if we took

this result literally, the peak fluorescence would be at zero energy (!). Probably the correct conclusion is that there is a tremendously strong coupling between excitation of the electrons and motion of the atoms, and presumably this is related to the fact that photon absorption leads to very rapid structural changes.

Before proceeding, it would be nice to do an honest calculation that reproduces the intuition of Figs 19 and 20, and this is done in Section A.3. The results of the calculation show, in more detail, how the coupling of electronic states to the vibrational motion of the molecule can shape the absorption spectrum. If there is just one lightly damped vibrational mode, then the single sharp absorption line which we expect from atomic physics becomes a sequence of lines, corresponding to changing electronic state and exciting one, two, three, ... or more vibrational quanta. If there are many modes, and these modes are damped by interaction with other degrees of freedom, these “vibronic” lines merge into a smooth spectrum which we can calculate in a semi-classical approximation.

The coupling of electronic transitions to vibrational motion generates the phenomenon of Raman scattering—a photon is inelastically scattered, making a virtual transition to the electronically excited state and dropping back down to the ground state, leaving behind a vibrational quantum **[add a figure illustrating Raman scattering]**. The energy shifts of the scattered photons allow us to read off, directly, the frequencies of the relevant vibrational modes. With a bit more sophistication, we can connect the strength of the different lines to the coupling constants (e.g., the displacements Δ_i along each mode, generalizing the discussion above) that characterize the interactions between electronic and vibrational degrees of freedom. If everything works, it should be possible to reconstruct the absorption spectrum from these estimates of frequencies and couplings. This whole program has been carried through for Rhodospin. Importantly, in order to get everything right, one has to include motions which are effectively unstable in the excited state, presumably corresponding to the torsional motions that lead to cis–trans isomerization. **[This is all a little quick. On the other hand, there is a huge amount of detail here that might take us away from the goal. Advice is welcome!]**

Problem 23: Raman scattering. Take the students through a simple calculation of Raman scattering ...

If we try to synthesize all of these ideas into a single schematic, we might get something like Fig 21. If we take




FIG. 21 Schematic model of the energy surfaces in Rhodopsin. The ground state has minima at both the 11-cis and the all-trans structures. A single excited state sits above this surface. At some intermediate structure, the surfaces come very close. At this point, the Born–Oppenheimer approximation breaks down, and there will be some mixing between the two states. A molecule lifted into the excited state by absorbing a photon slides down the upper surface, and can pass non-adiabatically into the potential well whose minimum is at all-trans.

this picture seriously, then after exciting the molecule with a pulse of light, we should see the disappearance of the absorption band associated with the 11-cis structure, the gradual appearance of the absorption from the all-trans state, and with a little luck, stimulated emission while the excited state is occupied. All of this is seen. Looking closely (e.g., at Fig 17), however, one sees that spectra are oscillating in time. Rather than sliding irreversibly down the potential surfaces toward their minima, the atomic structure oscillates. More remarkably, detailed analysis of the time evolution of the spectra demonstrates that there is coherent quantum mechanical mixing among the relevant electronic and vibrational states.

Our usual picture of molecules and their transitions comes from chemical kinetics: there are reaction rates, which represent the probability per unit time for the molecule to make transitions among states which are distinguishable by some large scale rearrangement; these transitions are cleanly separated from the time scales for molecules to come to equilibrium in each state. The initial isomerization event in rhodopsin is so fast that this approximation certainly breaks down. More profoundly, the time scale of the isomerization is so fast that it competes with the processes that destroy quantum mechanical coherence among the relevant electronic and vibrational states. The whole notion of an irreversible transition from one state to another necessitates the loss of coherence between these states (recall Schrödinger’s cat),

and so in this sense the isomerization is proceeding as rapidly as possible.

At this point what we would like to do is an honest, if simplified calculation that generates the schematic in Fig 21 and explains how the dynamics on these surfaces can be so fast. As far as I know, there is no clear answer to this challenge, although there are many detailed simulations, in the quantum chemical style, that probably capture elements of the truth.[\[it would be nice to be a little more explicit here!\]](#) The central ingredient is the special nature of the π bonds along the retinal. In the ground state, electron hopping between neighboring p_z orbitals lowers the energy of the system, and this effect is maximized in planar structures where the orbitals are all in the same orientation. But this lowering of the energy depends on the character of the electron wave functions—in the simplest case of bonding between two atoms, the symmetric state (the ‘bonding orbital’) has lower energy in proportion to the hopping matrix element, while the anti-symmetric state (‘anti-bonding orbital’) has higher energy, again in proportion to the matrix element. Thus, if we excite the electrons, it is plausible that the energy of the excited state could be reduced by structural changes that reduce the hopping between neighboring carbons, which happens if the molecule rotates to become non-planar. In this way we can understand why there is a force for rotation in the excited state, and why there is another local minimum in the ground state at the 11-cis structure.

Problem 24: Energy levels in conjugated molecules. The simplest model for a conjugated molecule is that the electrons which form the π orbitals can sit on each carbon atom with some energy that we can set to zero, and they can hop from one atom to its neighbors. Note that there is one relevant electron per carbon atom. If we write the Hamiltonian for the electrons as a matrix, then for a ring of six carbons (benzene) we have

$$\mathbf{H}_6 = \begin{pmatrix} 0 & -t & 0 & 0 & 0 & -t \\ -t & 0 & -t & 0 & 0 & 0 \\ 0 & -t & 0 & -t & 0 & 0 \\ 0 & 0 & -t & 0 & -t & 0 \\ 0 & 0 & 0 & -t & 0 & -t \\ -t & 0 & 0 & 0 & -t & 0 \end{pmatrix}, \quad (69)$$

where the “hopping matrix element” $-t$ is negative because the electrons can lower their energy by being shared among neighboring atoms—this is the essence of chemical bonding! Models like this are called *tight binding* models in the condensed matter physics literature and *Hückel* models in the chemical literature. Notice that they leave out any direct interactions among the electrons. This problem is about solving Schrödinger’s equation, $\mathbf{H}\psi = E\psi$, to find the energy eigenstates and the corresponding energy levels. Notice that for the case of benzene if we write the wave function ψ in terms of its six components (one for each carbon atom) then

Schrödinger's equation becomes

$$-t(\psi_2 + \psi_6) = E\psi_1 \quad (70)$$

$$-t(\psi_1 + \psi_3) = E\psi_2 \quad (71)$$

$$-t(\psi_2 + \psi_4) = E\psi_3 \quad (72)$$

$$-t(\psi_3 + \psi_5) = E\psi_4 \quad (73)$$

$$-t(\psi_4 + \psi_6) = E\psi_5 \quad (74)$$

$$-t(\psi_5 + \psi_1) = E\psi_6. \quad (75)$$

(a.) Considering first the case of benzene, show that solutions to the Schrödinger equation are of the form $\psi_n \propto \exp(ikn)$. What are the allowed values of the “momentum” k ? Generalize to an arbitrary N -membered ring.

(b.) What are the energies corresponding to the states labeled by k ? Because of the Pauli principle, the ground state of the molecule

is constructed by putting the electrons two-by-two (spin up and spin down) into the lowest energy states; thus the ground state of benzene has two electrons in each of the lowest three states. What is the ground state energy of benzene? What about for an arbitrary N -membered ring (with N even)? Can you explain why benzene is especially stable?

(c.) Suppose that the bonds between carbon atoms stretch and compress a bit, so that they become alternating single and double bonds rather than all being equivalent. To first order, if the bond stretches by an amount u then the hopping matrix element should go down (the electron has farther to hop), so we write $t \rightarrow t - \alpha u$; conversely, if the bond compresses, so that u is negative, the hopping matrix element gets larger. If we have alternating long and short (single and double) bonds, then the Hamiltonian for an six membered ring would be

$$\mathbf{H}_6(u) = \begin{pmatrix} 0 & -t + \alpha u & 0 & 0 & 0 & -t - \alpha u \\ -t + \alpha u & 0 & -t - \alpha u & 0 & 0 & 0 \\ 0 & -t - \alpha u & 0 & -t + \alpha u & 0 & 0 \\ 0 & 0 & -t + \alpha u & 0 & -t - \alpha u & 0 \\ 0 & 0 & 0 & -t - \alpha u & 0 & -t + \alpha u \\ -t - \alpha u & 0 & 0 & 0 & -t + \alpha u & 0 \end{pmatrix}. \quad (76)$$

Find the ground state energy of the electrons as a function of u , and generalize to the case of N -membered rings. Does the “dimerization” of the system ($u \neq 0$) raise or lower the energy of the electrons? Note that if your analytic skills (or patience!) give out, this is a relatively simple numerical problem; feel free to use the computer, but be careful to explain what units you are using when you plot your results.

(d.) In order to have bonds alternately stretched and compressed by an amount u , we need an energy $\frac{1}{2}\kappa u^2$ in each bond, where κ is the stiffness contributed by all the other electrons that we're not keeping track of explicitly. Consider parameter values $t = 2.5$ eV, $\alpha = 4.1$ eV/Å, and $\kappa = 21$ eV/Å². Should benzene have alternating single and double bonds ($u \neq 0$) or should all bonds be equivalent ($u = 0$)?

(e.) Peierls' theorem about one-dimensional electron systems predicts that, for N -carbon rings with N large, the minimum total energy will be at some non-zero u_* . Verify that this is true in this case, and estimate u_* . How large does N have to be before it's “large”? What do you expect for retinal?

I could try to do a full calculation here that puts flesh on the outline in the previous paragraph, using the tools from the problem above. But there still is a problem even if this works ...

Suppose that we succeed, and have a semi-quantitative theory of the excited state dynamics of rhodopsin, enough to understand why the quantum yield of fluorescence is so low, and what role is played by quantum coherence. We would then have to check that the barrier between the 11-cis and the all-trans structures in Fig 21 comes out to have the right height to explain the activation energy for spontaneous isomerization. But then how do we account for the anomalously low prefactor in

this rate, which is where, as discussed above, the protein acts to suppress dark noise? If there is something special about the situation in the environment of the protein which makes possible the ultrafast, coherent dynamics in the excited state, why does this special environment generate almost the same barrier as for isolated retinal?

It is clear that the ingredients for understanding the dynamics of rhodopsin—and hence for understanding why we can see into the darkest times of night—involve quantum mechanical ideas more related to condensed matter physics than to conventional biochemistry, a remarkably long distance from the psychology experiments on human subjects that we started with. While Lorentz could imagine that people count single quanta, surely he couldn't have imagined that he first steps of this process are coherent. While these are the ingredients, it is clear that we don't have them put together in quite the right way yet.

If rhodopsin were the only example of this “almost coherent chemistry” that would be good enough, but in fact the other large class of photon induced events in biological systems—photosynthesis—also proceed so rapidly as to compete with loss of coherence, and the crucial events again seem to happen (if you'll pardon the partisanship) while everything is still in the domain of physics and not conventional chemistry. Again there are beautiful experiments that present a number of theoretical challenges.²⁰

²⁰ As usual, a guide is found in the references at the end of this section.

Why biology pushes to these extremes is a good question. How it manages to do all this with big floppy molecules in water at roughly room temperature also is a great question.

To get some of the early history of work on the visual pigments, one can do worse than to read Wald's Nobel lecture (Wald 1972). Wald himself (along with his wife and collaborator, Ruth Hubbard) was quite an interesting fellow, much involved in politics; to connect with the previous section, his PhD adviser was Selig Hecht. [\[need more about dark noise and temperature dependence?\]](#) For a measurement of dark noise in cones, see Sampath & Baylor (2002). The remarkable result that the quantum yield of fluorescence in rhodopsin is $\sim 10^{-5}$ is due to Doukas et al (1984); it's worth noting that measuring this small quantum yield was possible at a time when one could not directly observe the ultrafast processes that are responsible for making the branching ratio this small. Direct measurements were finally made by Mathies et al (1988), Schoenlein et al (1991), and Wang et al (1994), the last paper making clear that the initial events are quantum mechanically coherent. A detailed analysis of the Raman spectra of Rhodopsin has been done by Loppnow & Mathies (1988).

Doukas et al 1984: Fluorescence quantum yield of visual pigments: Evidence for subpicosecond isomerization rates. AG Doukas, MR Junnarkar, RR Alfano, RH Callender, T Kakitani & B Honig, *Proc Nat Acad Sci (USA)* **81**, 4790–4794 (1984).

Loppnow & Mathies 1988: Excited-state structure and isomerization dynamics of the retinal chromophore in rhodopsin from resonance Raman intensities GR Loppnow & RA Mathies, *Biophys J* **54**, 35–43 (1988).

Mathies et al 1988: Direct observation of the femtosecond excited-state cis–trans isomerization in bacteriorhodopsin. RA Mathies, CH Brito Cruz, WT Pollard & CV Shank, *Science* **240**, 777–779 (1988).

Sampath & Baylor 2002: Molecular mechanisms of spontaneous pigment activation in retinal cones. AP Sampath & DA Baylor, *Biophys J* **83**, 184–193 (2002).

Schoenlein et al 1991: The first step in vision: Femtosecond isomerization of rhodopsin. RW Schoenlein, LA Peteanu, RA Mathies & CV Shank, *Science* **254**, 412–415 (1991).

Wald 1972: The molecular basis of visual excitation. G Wald, in *Nobel Lectures: Physiology or Medicine 1963–1970* (Elsevier, Amsterdam, 1972). Also available at <http://nobelprize.org>.

Wang et al 1994: Vibrationally coherent photochemistry in the femtosecond primary event of vision. Q Wang, RW Schoenlein, LA Peteanu, RA Mathies & CV Shank, *Science* **266**, 422–424 (1994).

The Born–Oppenheimer approximation is discussed in almost all quantum mechanics textbooks. For a collection of the key papers, with commentary, on the rich phenomena that can emerge in such adiabatic approximations, see Shapere & Wilczek (1989). Models for coupling of electron hopping to bond stretching (as in the last problem) were explored by Su, Schrieffer and Heeger in relation to polyacetylene. Importantly, these models predict that the excitations (e.g., upon photon absorption) are not just electrons and holes in the usual ladder of molecular orbitals, but that there are localized, mobile objects with unusual quantum numbers. These mobile objects can be generated by doping, which is the basis for conductivity in these quasi–one dimensional materials. The original work in Su et al (1980); a good review is Heeger et al (1988). Many people must have realized that the dynamical models being used by condensed matter physicists for (ideally) infinite chains

might also have something to say about finite chains. For ideas in this direction, including some specifically relevant to Rhodopsin, see Bialek et al (1987), Vos et al (1996), and Aalberts et al (2000).

Aalberts et al 2000: Quantum coherent dynamics of molecules: A simple scenario for ultrafast photoisomerization. DP Aalberts, MSL du Croo de Jongh, BF Gerke & W van Saarloos, *Phys Rev A* **61**, 040701 (2000).

Heeger et al 1988: Solitons in conducting polymers. AJ Heeger, S Kivelson, JR Schrieffer & W–P Su, *Rev Mod Phys* **60**, 781–850 (1988).

Bialek et al 1987: Simple models for the dynamics of biomolecules: How far can we go?. W Bialek, RF Goldstein & S Kivelson, in *Structure, Dynamics and Function of Biomolecules: The First EBSA Workshop*, A Ehrenberg, R Rigler, A Graslund & LJ Nilsson, eds, pp 65–69 (Springer–Verlag, Berlin, 1987).

Shapere & Wilczek 1989: *Geometric Phases in Physics* A Shapere and F Wilczek (World Scientific, Singapore, 1989)

Su et al 1980: Soliton excitations in polyacetylene. W–P Su, JR Schrieffer & AJ Heeger, *Phys Rev B* **22**, 2099–2111 (1980).

Vos et al 1996: Su–Schrieffer–Heeger model applied to chains of finite length. FLJ Vos, DP Aalberts & W van Saarloos, *Phys Rev B* **53**, 14922–14928 (1996).

Going beyond the case of rhodopsin, you may want to explore the role of quantum coherence in the initial events of photosynthesis; for an introduction see Fleming & van Grondelle (1994). The first experiments focused on photo–induced electron transfer, and looked at systems that had been genetically modified so that the electron, once excited, had no place to go (Vos et al 1991, Vos et al 1993); this made it possible to see the coherent vibrational motion of the molecule more clearly in spectroscopic experiments. Subsequent experiments used more intact systems, but looked first at low temperatures (Vos et al 1994a) and finally at room temperature (Vos et al 1994b). Eventually it was even possible to show that photo–triggering of electron transfer in other systems could reveal coherent vibrational motions (Liebl et al 1999). More or less at the same time as the original Vos et al experiments, my colleagues and I made the argument that photo–induced electron transfer rates in the initial events of photosynthesis would be maximized if the system were poised on the threshold of revealing coherent effects; maybe (although there were uncertainties about all the parameters) one could even strengthen this argument to claim that the observed rates were possible only in this regime (Skourtis et al 1992). Most recently, it has been discovered that when energy is trapped in the “antenna pigments” of photosynthetic systems, the migration of energy toward the reaction center (where the electron transfer occurs) is coherent, and it has been suggested that this allows for a more efficient exploration of space, finding the target faster than is possible in diffusive motion (Engel et al 2007). [\[Decide what to say about the large follow up literature!\]](#)

Engel et al 2007: Evidence for wavelike energy transfer through quantum coherence in photosynthetic systems. GS Engel, TR Calhoun, EL Read, T–K Ahn, T Mančal, Y–C Cheng, RE Blankenship & GR Fleming, *Nature* **446**, 782–786 (2007).

Fleming & van Grondelle 1994: The primary steps of photosynthesis. GR Fleming & R van Grondelle, *Physics Today* pp 48–55, February 1994.

Liebl et al 1999: Coherent reaction dynamics in a bacterial cytochrome c oxidase. U Liebl, G Lipowski, M Négrerie, JC Lambry, JL Martin & MH Vos, *Nature* **401**, 181–184 (1999).

Skourtis et al 1992: A new look at the primary charge separation in bacterial photosynthesis. SS Skourtis, AJR DaSilva, W Bialek & JN Onuchic, *J Phys Chem* **96**, 8034–8041 (1992).

Vos et al 1991: Direct observation of vibrational coherence in bacterial reaction centers using femtosecond absorption

spectroscopy. MH Vos, JC Lambry, SJ Robles, DC Youvan, J Breton & JL Martin, *Proc Nat'l Acad Sci (USA)* **88**, 8885–8889 (1991).

Vos et al 1993: Visualization of coherent nuclear motion in a membrane protein by femtosecond spectroscopy. MH Vos, F Rappaport, JC Lambry, J Breton & JL Martin, *Nature* **363**, 320–325 (1993).

Vos et al 1994a: Coherent dynamics during the primary electron transfer reaction in membrane-bound reaction centers of *Rhodobacter sphaeroides*. MH Vos, MR Jones, CN Hunter, J Breton, JC Lambry & JL Martin, *Biochemistry* **33**, 6750–6757 (1994).

Vos et al 1994b: Coherent nuclear dynamics at room temperature in bacterial reaction centers. MH Vos, MR Jones, CN Hunter, J Breton, JC Lambry & JL Martin, *Proc Nat'l Acad Sci (USA)* **91**, 12701–12705 (1994).

C. Dynamics of biochemical networks

Section still needs editing, as of September 18, 2011. The material here seems to have accreted during the early versions of the course, and much time is spent on things which we now know aren't productive On the other hand, I would like to say more about, for example, Sengupta et al (2000) on SNR in cascades and gain-bandwidth, as well as returning to the problem of transduction in invertebrates, e.g. theoretical work from Shraiman, Ranganathan et al.. So, I'd like make a more thorough overhaul here!

We have known for a long time that light is absorbed by rhodopsin, and that light absorption leads to an electrical response which is detectable as a modulation in the current flowing across the photoreceptor cell membrane. It is only relatively recently that we have come to understand the mechanisms which link these two events. The nature of the link is qualitatively different in different classes of organisms. For vertebrates, including us, the situation is as schematized in Fig 22. [it would be nice to come back and talk about invertebrates too]

In outline, what happens is that the excited rhodopsin changes its structure, arriving after several steps in a state where it can act as a catalyst to change the structure of another protein called transducin (T). The activated transducin in turn activates a catalyst called phosphodiesterase (PDE), which breaks down cyclic guanosine monophosphate (cGMP). Finally, cGMP binds to channels in the cell membrane and opens the channels, allowing current to flow (mostly carried by Na^+ ions); breaking down the cGMP thus decreases the number of open channels and decreases the current. [This discussion needs to refer to a schematic of the rod cell. Where is this? Earlier? Here?]

In a photomultiplier, photon absorption results in the ejection of a primary photoelectron, and then the large




FIG. 22 The cascade leading from photon absorption to ionic current flow in rod photoreceptors. Solid lines indicate ‘forward’ steps that generate gain; dashed lines are the ‘backward’ steps that shut off the process. **T** is the transducin molecule, a member of the broad class of G-proteins that couple receptors to enzymes. **PDE** is the enzyme phosphodiesterase, named for the particular bond that it cuts when it degrades cyclic guanosine monophosphate (**cGMP**) into **GMP**. **GC** is the guanylate cyclase that synthesizes **cGMP** from guanosine triphosphate, **GTP**.

electric field accelerates this electron so that when it hits the next metal plate it ejects many electrons, and the process repeats until at the output the number of electrons is sufficiently large that it constitutes a macroscopic current. Thus the photomultiplier really is an electron multiplier. In the same way, the photoreceptor acts as a molecule multiplier, so that for one excited rhodopsin molecule there are many cGMP molecules degraded at the output of the “enzymatic cascade.”

There are lots of interesting questions about how the molecule multiplication actually works in rod photoreceptors. These questions are made more interesting by the fact that this general scheme is ubiquitous in biological systems. [need a schematic about G-protein coupled receptors!] Rhodopsin is a member of a family of proteins which share common structural features (seven alpha helices that span the membrane in which the protein is embedded) and act as receptors, usually activated by the binding of small molecules such as hormones or odorants rather than light. Proteins in this family interact with proteins from another family, the G proteins, of which transducin is an example, and the result of such interactions typically is the activation of yet another enzyme, often one which synthesizes or degrades a cyclic nucleotide. Cyclic nucleotides in turn are common intracellular messengers, not just opening ion channels but also activating or inhibiting a variety of enzymes. This universality of components means that understanding the mechanisms of photon counting in rod cells is not just a curiosity for physicists, but a place where we can provide a model for understanding an enormous range of biological processes.

In order to get started, we need to know a little bit about ion channels, which form the output of the system. We will see that even the simplest, order-of-magnitude properties of channels raise a question about the observed behavior of the rod cells.

Recall that the brain contains no metallic or semiconductor components. Signals can still be carried by electrical currents and voltages, but now currents consist of ions, such as potassium or sodium, flowing through water or through specialized conducting pores. These pores, or channels, are large molecules (proteins) embedded in the cell membrane, and can thus respond to the electric field or voltage across the membrane as well as to the binding of small molecules. The coupled dynamics of channels and voltage turns each cell into a potentially complex nonlinear dynamical system.

Imagine a spherical molecule or ion of radius a ; a typical value for this radius is 0.3 nm. From Stokes' formula we know that if this ion moves through the water at velocity v it will experience a drag force $F = \gamma v$, with the drag coefficient $\gamma = 6\pi\eta a$, where η is the viscosity; for water $\eta = 0.01$ poise, the cgs unit poise = gm/(cm · s). The inverse of the drag coefficient is called the mobility, $\mu = 1/\gamma$, and the diffusion constant of a particle is related to the mobility and the absolute temperature by the Einstein relation or fluctuation dissipation theorem, $D = k_B T \mu$, with k_B being Boltzmann's constant and T the absolute temperature. Since life operates in a narrow range of absolute temperatures, it is useful to remember that at room temperature (25°C), $k_B T \sim 4 \times 10^{-21}$ J $\sim 1/40$ eV. So let's write the diffusion constant in terms of the other quantities, and then evaluate the order of magnitude:

$$D = k_B T \mu = k_B T \cdot \frac{1}{\gamma} = \frac{k_B T}{6\pi\eta a} \quad (77)$$

$$= \frac{[4 \times 10^{-21} \text{ J}]}{6\pi \cdot [0.01 \text{ gm}/(\text{cm} \cdot \text{s})] \cdot [0.3 \times 10^{-9} \text{ m}]} \quad (78)$$

$$\sim 2 \times 10^{-9} \text{ m}^2/\text{s} = 2 \mu\text{m}^2/\text{ms}. \quad (79)$$

Ions and small molecules diffuse freely through water, but cells are surrounded by a membrane that functions as a barrier to diffusion. In particular, these membranes are composed of lipids, which are nonpolar, and therefore cannot screen the charge of an ion that tries to pass through the membrane. The water, of course, is polar and does screen the charge, so pulling an ion out of the water and pushing it through the membrane would require surmounting a large electrostatic energy barrier. This barrier means that the membrane provides an enormous resistance to current flow between the inside and the outside of the cell. If this were the whole story there would be no electrical signaling in biology. In fact, cells construct specific pores or channels through which ions can pass, and by regulating the state of these channels the cell can control the flow of electric current across the membrane. [\[need a sketch that goes with this discussion\]](#)

Ion channels are themselves molecules, but very large ones—they are proteins composed of several thousand atoms in very complex arrangements. Let's try, however, to ask a simple question: If we open a pore in the cell membrane, how quickly can ions pass through? More precisely, since the ions carry current and will move in response to a voltage difference across the membrane, how large is the current in response to a given voltage?

Imagine that one ion channel serves, in effect, as a hole in the membrane. Let us pretend that ion flow through this hole is essentially the same as through water. The electrical current that flows through the channel is

$$J = q_{\text{ion}} \cdot [\text{ionic flux}] \cdot [\text{channel area}], \quad (80)$$

where q_{ion} is the charge on one ion, and we recall that 'flux' measures the rate at which particles cross a unit area, so that

$$\text{ionic flux} = \frac{\text{ions}}{\text{cm}^2 \text{s}} = \frac{\text{ions}}{\text{cm}^3} \cdot \frac{\text{cm}}{\text{s}} \quad (81)$$

$$= [\text{ionic concentration}] \cdot [\text{velocity of one ion}] \\ = cv. \quad (82)$$

Major current carriers such as sodium and potassium are at concentrations of $c \sim 100$ mM, or $c \sim 6 \times 10^{19}$ ions/cm³.

The next problem is to compute the typical velocity of one ion. We are interested in a current, so this is not the velocity of random Brownian motion but rather the average of that component of the velocity directed along the electric field. In a viscous medium, the average velocity is related to the applied force through the mobility, or the inverse of the drag coefficient as above. The force on an ion is in turn equal to the electric field times the ionic charge, and the electric field is (roughly) the voltage difference V across the membrane divided by the thickness ℓ of the membrane:

$$v = \mu F = \mu q_{\text{ion}} E \sim \mu q_{\text{ion}} \frac{V}{\ell} = \frac{D}{k_B T} q_{\text{ion}} \frac{V}{\ell}. \quad (83)$$

Putting the various factors together we find the current

$$J = q_{\text{ion}} \cdot [\text{ionic flux}] \cdot [\text{channel area}] \\ = q_{\text{ion}} \cdot [cv] \cdot [\pi d^2/4] \quad (84)$$

$$= q_{\text{ion}} \cdot \left[c \cdot \frac{D}{\ell} \cdot \frac{q_{\text{ion}} V}{k_B T} \right] \cdot \frac{\pi d^2}{4} \quad (85)$$

$$= \frac{\pi}{4} q_{\text{ion}} \cdot \frac{cd^2 D}{\ell} \cdot \frac{q_{\text{ion}} V}{k_B T}, \quad (86)$$

where the channel has a diameter d . If we assume that the ion carries one electronic charge, as does sodium, potassium, or chloride, then $q_{\text{ion}} = 1.6 \times 10^{-19}$ C and $q_{\text{ion}} V / (k_B T) = V / (25 \text{ mV})$. Typical values for the channel diameter should be comparable to the diameter of a single ion, $d \sim 0.3$ nm, and the thickness of the membrane is $\ell \sim 5$ nm. Thus

$$\begin{aligned}
J &= \frac{\pi}{4} q_{\text{ion}} \cdot \frac{cd^2D}{\ell} \cdot \frac{q_{\text{ion}}V}{k_B T} \\
&= \frac{\pi}{4} (1.6 \times 10^{-19} \text{ C}) \cdot \frac{(6 \times 10^{19} \text{ cm}^{-3})(3 \times 10^{-8} \text{ cm})^2(10^{-5} \text{ cm}^2/\text{s})}{50 \times 10^{-8} \text{ cm}} \cdot \frac{V}{25 \text{ mV}}
\end{aligned} \tag{87}$$

$$\sim 2 \times 10^{-14} \cdot \frac{V}{\text{mV}} \text{ C/s} \sim 2 \times 10^{-11} \frac{V}{\text{Volts}} \text{ Amperes}, \tag{88}$$

or

$$J = gV \tag{89}$$

$$g \sim 2 \times 10^{-11} \text{ Amperes/Volt} = 20 \text{ pS} \tag{90}$$

So our order of magnitude argument leads us to predict that the conductance of an open channel is roughly 20 pS.²¹ With a voltage difference across the membrane of ~ 50 mV, we thus expect that opening a single channel will cause ~ 1 picoAmp of current to flow. Although incredibly oversimplified, this is basically the right answer, as verified in experiments where one actually measures the currents flowing through single channel molecules.

The first problem in understanding the enzymatic cascade in rods is accessible just from these back of the envelope arguments. When we look at the total change in current that results from a single photon arrival, it is also ~ 1 pA. But if this were just the effect of (closing) one channel, we'd see "square edges" in the current trace as the single channels opened or closed. It would also be a little weird to have sophisticated (and expensive!) mechanisms for generating macroscopic changes in cGMP concentration only to have this act once again on a single molecule—if we have a single molecule input and a single molecule output, it really isn't clear why we would need an amplifier. What's going on?

The answer turns out to be that these channels flicker very rapidly between their open and closed states, so that on the relatively slow time scale of the rod response one sees essentially a graded current proportional to the probability of the channel being open. Thus the population of channels in the rod cell membrane produces a current that depends continuously on the concentration of cGMP. Alternatively, the noise variance that is associated with the random binary variable open/closed has been spread over a very broad bandwidth, so that in the frequency range of interest (recall that the single photon response is on a time scale of ~ 1 s) the noise is much reduced. This idea is made precise in the following problem, which you can think of as an introduction to the

analysis of noise in "chemical" systems where molecules fluctuate among multiple states.

Problem 25: Flickering channels. Imagine a channel that has two states, open and closed. There is a rate k_{open} at which the molecule makes transitions from the closed state to the open state, and conversely there is a rate k_{close} at which the open channels transition into the closed state. If we write the number of open channels as n_{open} , and similarly for the number of closed channels, this means that the deterministic kinetic equations are

$$\frac{dn_{\text{open}}}{dt} = k_{\text{open}}n_{\text{closed}} - k_{\text{close}}n_{\text{open}} \tag{91}$$

$$\frac{dn_{\text{close}}}{dt} = k_{\text{close}}n_{\text{open}} - k_{\text{open}}n_{\text{close}}, \tag{92}$$

or, since $n_{\text{open}} + n_{\text{closed}} = N$, the total number of channels,

$$\frac{dn_{\text{open}}}{dt} = k_{\text{open}}(N - n_{\text{open}}) - k_{\text{close}}n_{\text{open}} \tag{93}$$

$$= -(k_{\text{open}} + k_{\text{close}})n_{\text{open}} + k_{\text{open}}N. \tag{94}$$

For a single channel molecule, these kinetic equations should be interpreted as saying that an open channel has a probability $k_{\text{close}}dt$ of making a transition to the closed state within a small time dt , and conversely a closed channel has a probability $k_{\text{open}}dt$ of making a transition to the open state. We will give a fuller account of noise in chemical systems in the next Chapter, but for now you should explore this simplest of examples.

(a.) If we have a finite number of channels, then really the number of channels which make the transition from the closed state to the open state in a small window dt is a random number. What is the mean number of these closed \rightarrow open transitions? What is the mean number of open \rightarrow closed transitions? Use your results to show that macroscopic kinetic equations such as Eqs (91) and (92) should be understood as equations for the mean numbers of open and closed channels,

$$\frac{d\langle n_{\text{open}} \rangle}{dt} = k_{\text{open}}\langle n_{\text{closed}} \rangle - k_{\text{close}}\langle n_{\text{open}} \rangle \tag{95}$$

$$\frac{d\langle n_{\text{close}} \rangle}{dt} = k_{\text{close}}\langle n_{\text{open}} \rangle - k_{\text{open}}\langle n_{\text{close}} \rangle. \tag{96}$$

(b.) Assuming that all the channels make their transitions independently, what is the variance in the number of closed \rightarrow open transitions in the small window dt ? In the number of open \rightarrow closed transitions? Are these fluctuations in the number of transitions independent of one another?

(c.) Show that your results in [b] can be summarized by saying that the change in the number of open channels during the time dt obeys an equation

$$n_{\text{open}}(t+dt) - n_{\text{open}}(t) = dt[k_{\text{open}}n_{\text{closed}} - k_{\text{close}}n_{\text{open}}] + \eta(t), \tag{97}$$

where $\eta(t)$ is a random number that has zero mean and a variance

$$\langle \eta^2(t) \rangle = dt[k_{\text{open}}n_{\text{closed}} + k_{\text{close}}n_{\text{open}}]. \tag{98}$$

²¹ Siemens are the units of conductance, which are inverse to units of resistance, ohms. In the old days, this inverse of resistance had the rather cute unit 'mho' (pronounced 'moe,' like the Stooge).

Explain why the values of $\eta(t)$ and $\eta(t')$ are independent if $t \neq t'$.

(d.) This discussion should remind you of the description of Brownian motion by a Langevin equation, in which the deterministic dynamics are supplemented by a random force that describes molecular collisions. In this spirit, show that, in the limit $dt \rightarrow 0$, you can rewrite your results in [c] to give a Langevin equation for the number of open channels,

$$\frac{dn_{\text{open}}}{dt} = -(k_{\text{open}} + k_{\text{close}})n_{\text{open}} + k_{\text{open}}N + \zeta(t), \quad (99)$$

where

$$\langle \zeta(t)\zeta(t') \rangle = \delta(t-t')[k_{\text{open}}n_{\text{closed}} + k_{\text{close}}n_{\text{open}}]. \quad (100)$$

In particular, if the noise is small, show that $n_{\text{open}} = \langle n_{\text{open}} \rangle + \delta n_{\text{open}}$, where

$$\frac{d\delta n_{\text{open}}}{dt} = -(k_{\text{open}} + k_{\text{close}})\delta n_{\text{open}} + \zeta_s(t), \quad (101)$$

$$\langle \zeta_s(t)\zeta_s(t') \rangle = 2k_{\text{open}}\langle n_{\text{closed}} \rangle. \quad (102)$$

(e.) Solve Eq (101) to show that

$$\langle \delta n_{\text{open}}^2 \rangle = Np_{\text{open}}(1 - p_{\text{open}}) \quad (103)$$

$$\langle \delta n_{\text{open}}(t)\delta n_{\text{open}}(t') \rangle = \langle \delta n_{\text{open}}^2 \rangle \exp\left[-\frac{|t-t'|}{\tau_c}\right], \quad (104)$$

where the probability of a channel being open is $p_{\text{open}} = k_{\text{open}}/(k_{\text{open}} + k_{\text{close}})$, and the correlation time $\tau_c = 1/(k_{\text{open}} + k_{\text{close}})$. Explain how the result for the variance $\langle \delta n_{\text{open}}^2 \rangle$ could be derived more directly.

(f.) Give a critical discussion of the approximations involved in writing down these Langevin equations. In particular, in the case of Brownian motion of a particle subject to ordinary viscous drag, the Langevin force has a Gaussian distribution. Is that true here?

Problem 26: Averaging out the noise. Consider a random variable such as n_{open} in the previous problem, for which the noise has exponentially decaying correlations, as in Eq (104). Imagine that we average over a window of duration τ_{avg} , to form a new variable

$$z(t) = \frac{1}{\tau_{\text{avg}}} \int_0^{\tau_{\text{avg}}} d\tau \delta n_{\text{open}}(t - \tau). \quad (105)$$

Show that, for $\tau_{\text{avg}} \gg \tau_c$, the variance of z is smaller than the variance of δn_{open} by a factor of τ_{avg}/τ_c . Give some intuition for why this is true (e.g., how many statistically independent samples of n_{open} will you see during the averaging time?). What happens if your averaging time is shorter?

I think this is a fascinating example, because evolution has selected for very fast channels to be present in a cell that signals very slowly! Our genome (as well as those of many other animals) codes for hundreds if not thousands of different types of channels once one includes the possibility of alternative splicing. These different channels differ, among other things, in their kinetics. In the fly retina, for example, the dynamics of visual inputs looking straight ahead are very different from those looking to the side, and in fact the receptor cells that look in these different directions have different kinds of channels—the faster channels to respond to the more rapidly varying signals. [I am not sure that the last statement is correct, and need to check the references; what certainly is true is that insects with different lifestyles (e.g., acrobats vs. slow fliers) use different potassium channels ...] In the




FIG. 23 Current through the rod cell membrane as a function of the cyclic GMP concentration. The fit is to Eq (106), with $n = 2.9 \pm 0.1$ and $G_{1/2} = 45 \pm 4 \mu\text{M}$. From Rieke & Baylor (1996).

vertebrate rod, signals are very slow but the channels are fast, and this makes sense only if the goal is to suppress the noise.

Having understood a bit about the channels, let's take one step back and see how these channels respond to cyclic GMP. Experimentally, with the rod outer segment sucked into the pipette for measuring current, one can break off the bottom of the cell and make contact with its interior, so that concentrations of small molecules inside the cell will equilibrate with concentrations in the surrounding solution. Since the cell makes cGMP from GTP, if we remove GTP from the solution then there is no source other than the one that we provide, and now we can map current vs concentration. The results of such an experiment are shown in Fig 23. We see that the current I depends on the cGMP concentration G as

$$I = I_{\text{max}} \frac{G^n}{G^n + G_{1/2}^n}, \quad (106)$$

with $n \approx 3$. This suggests strongly that the channel opens when three molecules of cGMP bind to it. This is an example of “cooperativity” or “allostery,” which is a very important theme in biochemical signaling and regulation. It's a little off to the side of our discussion here, however, so see Appendix A.4.

Let's try to write a more explicit model for the dynamics of amplification in the rod cell, working back from the channels. We have Eq (106), which tells us how the current I depends on G , the concentration of cyclic GMP. The dynamics of G has two terms, synthesis and degradation:

$$\frac{dG}{dt} = \gamma - PDE^*G, \quad (107)$$

where γ denotes the rate of synthesis by the guanylate cyclase (GC, cf Fig 22), and PDE^* measures the activity of the active phosphodiesterase. It turns out that there

is a feedback mechanism in the rod, where calcium enters through the open channels (as part of the current), and then calcium binding inhibits the activity of the guanylate cyclase. We can summarize these effects, measured in several experiments, by writing

$$\gamma = \frac{\gamma_{\max}}{1 + (Ca/K_{\text{gc}})^2} \approx \alpha Ca^{-2}, \quad (108)$$

where the last approximation is valid so long as the typical calcium concentration Ca is much larger than the binding constant $K_{\text{gc}} \sim 100$ nM, which seems to be true; the fact that the dependence is on the square of the calcium concentration presumably means that two Ca^{++} ions bind to inhibit the cyclase (see again the discussion of cooperativity in Appendix A.4). Since calcium enters the cell as a fraction of the current flowing through the open channels, and presumably is pumped back out by other mechanisms, we can write

$$\frac{dCa}{dt} = fI(G) - \beta Ca, \quad (109)$$

where f is the fraction of the current carried by calcium and $1/\beta$ is the lifetime of calcium before it is pumped out. These equations tell how the cyclic GMP concentration, and hence the current, will respond to changes in the activity of the phosphodiesterase, thus describing the last steps of the amplification cascade.

It is convenient to express the response of G to PDE^* in the limit that the response is linear, which we expect is right when only small numbers of photons are being counted. This linearization gives us

$$\delta\dot{G} = \frac{\partial\gamma}{\partial Ca} \delta Ca - PDE_0^* \delta G - G_0 \delta PDE^* \quad (110)$$

$$\delta\dot{Ca} = fI'(G_0) \delta G - \beta \delta Ca, \quad (111)$$

where the subscript 0 denotes the values in the dark. We can solve these equations by passing to Fourier space, where

$$\delta\tilde{G}(\omega) = \int dt e^{+i\omega t} \delta G(t), \quad (112)$$

and similarly for the other variables. As usual, this reduces the linear differential equations to linear algebraic equations, and when the dust settles we find

$$\frac{\delta\tilde{G}(\omega)}{\delta P\tilde{D}E^*(\omega)} = \frac{-G_0(-i\omega + \beta)}{(-i\omega + PDE_0^*)(-i\omega + \beta) + A}, \quad (113)$$

$$A = 2\gamma_0 f I'(G_0) / Ca_0. \quad (114)$$

Already this looks like lots of parameters, so we should see how we can simplify, or else measure some of the parameters directly.

First, one finds experimentally that the cyclic GMP concentration is in the regime where $I \propto G^3$, that is

$G \ll G_{1/2}$. This means that we can express the response more compactly as a fractional change in current

$$\delta\tilde{I}(\omega) = 3I_0 \frac{-i\omega + \beta}{(-i\omega + PDE_0^*)(-i\omega + \beta) + A} \cdot \delta P\tilde{D}E^*(\omega), \quad (115)$$

where $A = 6\beta PDE_0^*$.

Problem 27: Dynamics of cGMP. Fill in all the steps leading to Eq (115).

In the same experiment where one measures the response of the channels to cGMP, one can suddenly bring the cGMP concentration of the outside solution to zero, and then the internal cGMP concentration (which we can read off from the current, after the first experiment) will fall due both to diffusion out of the cell and to any PDE which is active in the dark; one can also poison the PDE with a drug (IBMX), separating the two components. In this way one can measure $PDE_0^* = 0.1 \pm 0.02 \text{ s}^{-1}$. To measure β , you need to know that the dominant mechanism for pumping calcium out of the cell actually generates an electrical current across the membrane.²² With this knowledge, if we turn on a bright light and close all the cGMP-sensitive channels, there is no path for calcium to enter the rod outer segment, but we still see a small current as it is pumped out. This current decays with a rate $\beta \sim 2 \text{ s}^{-1}$. Thus, although this model—even for part of the process!—looks complicated, there are many independent experiments one can do to measure the relevant parameters.

In fact, the analysis of the dynamics of cGMP and calcium leads us to the point where we can more or less invert these dynamics, turning the dynamics of the current back into the dynamics of the PDE^* . An interesting application of this idea is to try and understand the continuous background noise that occurs in the dark. As we saw, there is a big source of noise in the dark that comes from spontaneous isomerization of rhodopsin. But there is also a smaller, continuous rumbling, with an amplitude $\delta I_{\text{rms}} \sim 0.1$ pA. This isn't the intrinsically random opening and closing of the channels, since we have seen that this happens very fast and thus contributes very little to the noise at reasonable frequencies. It must thus reflect

²² This needn't be true. First, there are mechanisms which exchange ions on different sides of the membrane, maintaining electrical neutrality. Second, it could be that the dominant pump sends calcium into storage spaces inside the cell, so no ions cross the cell membrane.

responses of the channels to fluctuations in the concentration of cGMP. Since this concentration is determined by a balance between synthesis and degradation, one should check whether one of these processes is dominating the noise.

The rate at which cGMP is synthesized is modulated by calcium, but we can prevent the calcium concentration from changing by using buffers, either injected into the cell or in the surrounding solution when the cell is broken open. If the calcium concentration were itself fluctuating, and these fluctuations generated noise in the synthesis of cGMP, buffering the calcium concentration should lower the continuous background noise; instead the noise goes up. On the other hand, if we poison the phosphodiesterase with IBMX, and allow synthesis to compete with diffusion out of a broken cell, the noise drops dramatically. [At this point things get a little vague .. go back and do better!] These, and other experiments as well, indicate that the dominant source of the continuous dark noise is fluctuations in the number of active phosphodiesterase molecules. Alternatively, one can say that the noise arises from ‘spontaneous’ activation of the PDE, absent any input from activated rhodopsin.

[Need to be sure we have control over the math here .. maybe connect back to problem about ion channels? Also connect to Appendix A.2. Review before giving results. Get all the number right, too!] If the activation of PDE in the dark is rare, then we expect that the variance in the number of active molecules will be equal to the mean, and the fluctuations in activity should have a correlation time equal to the lifetime of the activated state. If a is the activity of a single enzyme—that is, the factor that converts the number of active enzymes into the rate at which cGMP is degraded—then we have

$$\langle \delta PDE^*(t) \delta PDE^*(t') \rangle = a PDE_0^* e^{-|t-t'|/\tau_c}, \quad (116)$$

where τ_c is the lifetime of the active state. Putting this together with Eq (115), we can generate a prediction for the power spectrum of fluctuations in the current. Importantly, the only unknown parameters are a , which sets the over scale of the fluctuations, and τ_c , which shapes the spectrum. Fitting to the observed spectra, one finds $a = 1.6 \times 10^{-5} \text{ s}^{-1}$ and $\tau_c = 0.56 \text{ s}$. Thus, a single active phosphodiesterase causes the cGMP concentration to decrease at a rate $aG_0 \sim 2 \times 10^{-4} \mu\text{M/s}$, and this lasts for roughly half a second; with a volume of $\sim 10^{-12} \text{ l}$, this means that one PDE^* destroys ~ 60 molecules of cGMP.

Knowing how changes in concentration change the current, and how much one PDE^* can reduce the cGMP concentration, we can calculate that a single photon must activate at least 2000 phosphodiesterase molecules. More concretely, a single activated rhodopsin must trigger the activation of at least 2000 PDE^* . In order for this to happen, the activated rhodopsin has to diffuse in the disk membrane [did we actually discuss the geometry of the disk etc? check!] during its lifetime; certainly the

number of molecules that it can activate is limited by the number of molecules that it can encounter via diffusion. With measured diffusion constants and a lifetime of roughly one second (after this, the whole response starts to shut off), this seems possible, but not with much to spare. Thus, it seems likely that the gain in the first part of the amplifier is limited by the density of molecules and the physics of diffusion. [Need estimates of diffusion constant here .. either explain, or give problem, about diffusion limit to this reaction.]

[I think that before going on to discuss reproducibility we want to say a bit more about gain .. look at Detwiler et al (2000) regarding the design of G protein elements, since this would also give an excuse to discuss some more about these ... Then check segue.] So, given this dissection of the amplifier, what is it that we really want to know? Understanding gain—how you get many molecules out for only one molecule at the input—isn’t so hard, basically because catalysis rates are high, close to the diffusion limit. One might want to understand the system’s choice of other parameters, but is there really a conceptual problem here?

Perhaps the most surprising aspect of the single photon response in rods is its reproducibility. If we look at the responses to dim light flashes and isolate those responses that correspond to a single photon (you have already done a problem to assess how easy or hard this is!), one finds that the amplitude of the response fluctuates by only $\sim 15 - 20\%$; see, for example, Fig. 24. To understand *why* this is surprising we have to think about chemistry at the level of single molecules, specifically the chemical reactions catalyzed by the single activated molecule of rhodopsin.

[This discussion need to point back to the problem about ion channels.] When we write that there is a rate k for a chemical reaction, what we mean is that for one molecule there is a *probability per unit time* k that the reaction will occur—this should be familiar from the case of radioactive decay. Thus when one molecule of rhodopsin is activated at time $t = 0$, if we imagine that de-activation is a simple chemical reaction then the probability that the molecule is still active at time t obeys the usual kinetic equation

$$\frac{dp(t)}{dt} = -kp(t); \quad (117)$$

of course if there are N total molecules then $Np(t) = n(t)$ is the expected number of molecules still in the active state. Thus, $p(t) = \exp(-kt)$. The probability density $P(t)$ that the molecule is active for exactly a time t is the probability that the molecule is still active at t times the probability per unit time of de-activation, so

$$P(t) = kp(t) = k \exp(-kt). \quad (118)$$

This may seem pedantic, but it’s important to be clear—and we’ll see that far from being obvious there must be




FIG. 24 Reproducibility of the single photon response, from Field & Rieke (2002b). (A) Examples of single photon responses and failures from single mammalian rods. (B) Variances of the responses in (A). (C) Variance and square of mean response to one photon; variance in the response is defined as the difference in variance between responses and failures. Finally (D) shows the mean of results as in (C) from eight primate rods and nine guinea pig rods; scales are normalized for each cell by the peak mean response and the time to peak. We see that at the peak response the relative variance is ~ 0.025 , so the root-mean-square fluctuations are ~ 0.15 .

something wrong with this simple picture.

Given the probability density $P(t)$, we can calculate the mean and variance of the time spent in the active state:

$$\langle t \rangle \equiv \int_0^\infty dt P(t) t \quad (119)$$

$$= k \int_0^\infty \exp(-kt) t = 1/k; \quad (120)$$

$$\langle (\delta t)^2 \rangle \equiv \int_0^\infty dt P(t) t^2 - \langle t \rangle^2 \quad (121)$$

$$= k \int_0^\infty dt \exp(-kt) t^2 - 1/k^2 \quad (122)$$

$$= 2/k^2 - 1/k^2 = 1/k^2. \quad (123)$$

Thus we find that

$$\delta t_{\text{rms}} \equiv \sqrt{\langle (\delta t)^2 \rangle} = 1/k = \langle t \rangle, \quad (124)$$

so that the root-mean-square fluctuations in the lifetime are equal to the mean.

How does this relate to the reproducibility of the single photon response? The photoreceptor works by having the active rhodopsin molecule act as a catalyst, activating transducin molecules. If the catalysis proceeds at some constant rate (presumably set by the time required for rhodopsin and transducin to find each by diffusion in the membrane), then the number of activated transducins is proportional to the time that rhodopsin spends in the active state—and hence we would expect that the number of active transducin molecules has root-mean-square fluctuations equal to the mean number. If the subsequent events in the enzymatic cascade again have outputs proportional to their input number of molecules, this variability will not be reduced, and the final output (the change in cGMP concentration) will again have relative fluctuations of order one, much larger than the observed 15–20%. This is a factor of 25 or 40 error in variance; we can't even claim to have an order of magnitude understanding of the reproducibility. I'd like to give an idea of the different possible solutions that people have considered, focusing on very simple versions of these ideas that we can explore analytically. At the end, we'll look at the state of the relevant experiments.

One possibility is that although the lifetime of activated rhodopsin might fluctuate, the number of molecules at the output of the cascade fluctuates less because of saturation [point to sketch of discs]. For example, if each rhodopsin has access only to a limited pool of transducin molecules, a reasonable fraction of rhodopsins might remain active long enough to hit all the molecules in the pool. The simplest version of this idea is as follows. Let the total number of transducins in the pool be N_{pool} , and let the number of activated transducins be n_T . When the rhodopsin is active, it catalyzes the conversion of inactive transducins (of which there are $N_{\text{pool}} - n_T$) into the active form at a rate r , so that (neglecting the discreteness of the molecules)

$$\frac{dn_T}{dt} = r(N_{\text{pool}} - n_T). \quad (125)$$

If the rhodopsin molecule is active for a time t then this catalysis runs for a time t and the number of activated transducins will be

$$n_T(t) = N_{\text{pool}}[1 - \exp(-rt)]. \quad (126)$$

For small t the variations in t are converted into proportionately large variations in n_T , but for large t the saturation essentially cuts off this variation.

To be more precise, recall that we can find the distribution of n_T by using the identity

$$P(n_T)dn_T = P(t)dt, \quad (127)$$

which applies whenever we have two variables that are related by a deterministic, invertible transformation. From

Eq (126) we have

$$t = -\frac{1}{r} \ln(1 - n_T/N_{\text{pool}}), \quad (128)$$

and so, going through the steps explicitly:

$$P(n_T) = P(t) \left| \frac{dn_T}{dt} \right|^{-1} \quad (129)$$

$$= k \exp(-kt) \frac{1}{r(N_{\text{pool}} - n_T)} \quad (130)$$

$$= \frac{k}{r} \exp \left[\left(\frac{k}{r} \right) \ln(1 - n_T/N_{\text{pool}}) \right] \frac{1}{(N_{\text{pool}} - n_T)} \quad (131)$$

$$= \frac{k}{rN_{\text{pool}}} \left(1 - \frac{n_T}{N_{\text{pool}}} \right)^{k/r-1}. \quad (132)$$

[Maybe a plot to show this?] When the activation rate r is small, n_T always stays much less than N_{pool} and the power law can be approximated as an exponential. When r is large, however, the probability distribution grows a power law singularity at N_{pool} ; for r finite this singularity is integrable but as $r \rightarrow \infty$ it approaches a log divergence, which means that essentially all of the weight will be concentrated at N_{pool} . In particular, the relative variance of n_T vanishes as r becomes large, as promised.

This discussion has assumed that the limited number of target molecules is set, perhaps by some fixed structural domain. Depending on details, it is possible for such a limit to arise dynamically, as a competition between diffusion and chemical reactions. In invertebrate photoreceptors, such as the flies we have met in our discussion above, there is actually a positive feedback loop in the amplifier which serves to ensure that each structural domain (which are more obvious in the fly receptor cells) ‘fires’ a saturated, stereotyped pulse in response to each photon.

[Make a sketch of the different models—either one big figure, or separate ones for each model.]

The next class of models are those that use feedback. The idea, again, is simple: If the output of the cascade is variable because the rhodopsin molecule doesn’t “know” when to de-activate, why not link the de-activation to the output of the cascade? Roughly speaking, count the molecules at the output and shut the rhodopsin molecule off when we reach some fixed count. Again let’s try the simplest version of this. When rhodopsin is active it catalyzes the formation of some molecule (which might not actually be the transducin molecule itself) at rate r , and let the number of these output molecules be x so that we simply have

$$\frac{dx}{dt} = r, \quad (133)$$

or $x = rt$. Let’s have the rate of deactivation of rhodopsin depend on x , so that instead of Eq (117) we have

$$\frac{dp(t)}{dt} = k[x(t)]p(t). \quad (134)$$

For example, if deactivation is triggered by the cooperative binding of m x molecules (as in the discussion of cGMP-gated channels), we expect that

$$k[x] = k_{\text{max}} \frac{x^m}{x_0^m + x^m}. \quad (135)$$

We can solve Eq (134) and then recover the probability density for rhodopsin lifetime as before,

$$p(t) = \exp \left(- \int_0^t d\tau k[x(\tau)] \right) \quad (136)$$

$$P(t) = k[x(t)] \exp \left(- \int_0^t d\tau k[x(\tau)] \right). \quad (137)$$

Again we can push through the steps:

$$P(t) = k[x(t)] \exp \left(- \int_0^t d\tau k[x(\tau)] \right) = k_{\text{max}} \frac{x^m(t)}{x_0^m + x^m(t)} \exp \left(-k_{\text{max}} \int_0^t d\tau \frac{x^m(\tau)}{x_0^m + x^m(\tau)} \right) \quad (138)$$

$$\approx k_{\text{max}} \left(\frac{t}{t_0} \right)^m \exp \left[-\frac{k_{\text{max}} t_0}{m+1} \left(\frac{t}{t_0} \right)^{m+1} \right], \quad (139)$$

where in the last step we identify $t_0 = x_0/r$ and assume that $t \ll t_0$.

To get a better feel for the probability distribution in

Eq (139) it is useful to rewrite it as

$$P(t) \approx k_{\text{max}} \exp [-G(t)] \quad (140)$$

$$G(t) = -m \ln \left(\frac{t}{t_0} \right) + \frac{k_{\text{max}} t_0}{m+1} \left(\frac{t}{t_0} \right)^{m+1} \quad (141)$$

We can find the most likely value of the lifetime, \bar{t} , by minimizing G , which of course means that the derivative must be set to zero:

$$G'(t) = -\frac{m}{t} + k_{\max}t_0 \cdot \frac{1}{t} \left(\frac{t}{t_0}\right)^{m+1} \quad (142)$$

$$G'(t = \bar{t}) = 0 \Rightarrow k_{\max}t_0 \cdot \frac{1}{\bar{t}} \left(\frac{\bar{t}}{t_0}\right)^{m+1} = \frac{m}{\bar{t}} \quad (143)$$

$$\frac{\bar{t}}{t_0} = \left(\frac{m}{k_{\max}t_0}\right)^{1/m} \quad (144)$$

In particular we see that for sufficiently large k_{\max} we will have $\bar{t} \ll t_0$, consistent with the approximation above. What we really want to know is how sharp the distribution is in the neighborhood of \bar{t} , so we will try a series expansion of $G(t)$:

$$P(t) \approx k_{\max} \exp \left[-G(\bar{t}) - \frac{1}{2}G''(\bar{t})(t - \bar{t})^2 - \dots \right] \quad (145)$$

$$G''(t) = \frac{m}{t^2} + (k_{\max}t_0)m \cdot \frac{1}{t^2} \left(\frac{t}{t_0}\right)^{m+1} \approx \frac{m}{\bar{t}^2}, \quad (146)$$

where again in the last step we assume $\bar{t} \ll t_0$. Thus we see that the distribution of lifetimes is, at least near its peak,

$$P(t) \approx P(\bar{t}) \exp \left[-\frac{m}{2\bar{t}^2}(t - \bar{t})^2 - \dots \right]. \quad (147)$$

This of course is a Gaussian with variance

$$\langle(\delta t)^2\rangle = \frac{1}{m} \cdot \bar{t}^2, \quad (148)$$

so the relative variance is $1/m$ as opposed to 1 in the original exponential distribution.

A concrete realization of the feedback idea can be built around the fact that the current flowing into the rod includes calcium ions, and the resulting changes in calcium concentration can regulate protein kinases—proteins which in turn catalyze the attachment of phosphate groups to other proteins—and rhodopsin shut off is known to be associated with phosphorylation at multiple sites. Calcium activation of kinases typically is cooperative, so $m \sim 4$ in the model above is plausible. Notice that in the saturation model the distribution of lifetimes remains broad and the response to these variations is truncated; in the feedback model the distribution of lifetimes itself is sharpened.

A third possible model involves multiple steps in rhodopsin de-activation. Let us imagine that rhodopsin starts in one state and makes a transition to state 2, then from state 2 to state three, and so on for K states, and then it is the transition from state K to $K + 1$ that actually corresponds to de-activation. Thus there are K active states and if the time spent in each state is t_i then the total time spent in activated states is

$$t = \sum_{i=1}^K t_i. \quad (149)$$

Clearly the mean value of t is just the sum of the means of each t_i , and if the transitions are independent (again, this is what you mean when you write the chemical kinetics with the arrows and rate constants) then the variance of t will also be the sum of the variances of the individual t_i ,

$$\langle t \rangle = \sum_{i=1}^K \langle t_i \rangle \quad (150)$$

$$\langle(\delta t)^2\rangle = \sum_{i=1}^K \langle(\delta t_i)^2\rangle. \quad (151)$$

We recall from above that for each single step, $\langle(\delta t_i)^2\rangle = \langle t_i \rangle^2$. If the multiple steps occur at approximately equal rates, we can write

$$\langle t \rangle = \sum_{i=1}^K \langle t_i \rangle \approx K \langle t_1 \rangle \quad (152)$$

$$\langle(\delta t)^2\rangle = \sum_{i=1}^K \langle(\delta t_i)^2\rangle = \sum_{i=1}^K \langle t_i \rangle^2 \approx K \langle t_1 \rangle^2 \quad (153)$$

$$\frac{\langle(\delta t)^2\rangle}{\langle t \rangle^2} \approx \frac{K \langle t_1 \rangle^2}{(K \langle t_1 \rangle)^2} = \frac{1}{K}. \quad (154)$$


Thus the relative variance declines as one over the number of steps, and the relative standard deviation declines as one over the square root of the number of steps. This is an example of how averaging K independent events causes a $1/\sqrt{K}$ reduction in the noise level.

The good news is that allowing de-activation to proceed via multiple steps can reduce the variance in the lifetime of activated rhodopsin. Again our attention is drawn to the fact that rhodopsin shut off involves phosphorylation of the protein at multiple sites. The bad news is that to have a relative standard deviation of $\sim 20\%$ would require 25 steps.

It should be clear that a multistep scenario works only if the steps are irreversible. If there are significant “backward” rates then progress through the multiple states becomes more like a random walk, with an accompanying increase in variance. Thus each of the (many) steps involved in rhodopsin shut off must involve dissipation of a few $k_B T$ of energy to drive the whole process forward.

Problem 28: Getting the most out of multiple steps.

Consider the possibility that Rhodopsin leaves its active state through a two step process. To fix the notation, let's say that the first step occurs with a rate k_1 and the second occurs with rate k_2 :



Assume that we are looking at one molecule, and at time $t = 0$ this molecule is in state Rh^* .

(a) Write out and solve the differential equations for the time dependent probability of being in each of the three states.

(b) Use your results in [a] to calculate the probability distribution for the time at which the molecule *enters* the inactive state. This is the distribution of “lifetimes” for the two active states. Compute the mean and variance of this lifetime as a function of the parameters k_1 and k_2 .

(c) Is there a simple, intuitive argument that allows you to write down the mean and variance of the lifetime, without solving any differential equations? Can you generalize this to a scheme in which inactivation involves N steps rather than two?

(d) Given some desired mean lifetime, is there a way of adjusting the parameters k_1 and k_2 (or, more generally, $k_1, k_2 \dots, k_N$) to minimize the variance?

(e) Suppose that there is a back reaction $\text{Rh}^{**} \xrightarrow{k_{-1}} \text{Rh}^*$. Discuss what this does to the distribution of lifetimes. In particular, what happens if the rate k_{-1} is very fast? Note that “discuss” is deliberately ambiguous; you could try to solve the relevant differential equations, or to intuit the answer, or even do a small simulation [connect this problem to recent work by Escola & Paninski].

The need for energy dissipation and the apparently very large number of steps suggests a different physical picture. If there really are something like 25 steps, then if we plot the free energy of the rhodopsin molecule as a function of its atomic coordinates, there is a path from initial to final state that passes over 25 hills and valleys. Each valley must be a few $k_B T$ lower than the last, and the hills must be many $k_B T$ high to keep the rates in the right range. This means that the energy surface is quite rough [this needs a sketch]. Now when we take one solid and slide it over another, the energy surface is rough on the scale of atoms because in certain positions the atoms on each surface “fit” into the interatomic spaces on the other surface, and then as we move by an Ångström or so we encounter a very high barrier. If we step back and blur our vision a little bit, all of this detailed roughness just becomes friction between the two surfaces. Formally, if we think about Brownian motion on a rough energy landscape and we average over details on short length and time scales, what we will find is that the mobility or friction coefficient is renormalized and then the systems behaves on long time scales as if it were moving with this higher friction on a smooth surface.

So if the de-activation of rhodopsin is like motion on a rough energy surface, maybe we can think about the renormalized picture of motion on a smooth surface with high drag or low mobility. Suppose that the active and inactive states are separated by a distance ℓ along some direction in the space of molecular structures, and that motion in this direction occurs with an effective mobility μ . If there is an energy drop ΔE between the active and de-activated states, then the velocity of motion is $v \sim \mu \Delta E / \ell$ and the mean time to make the de-activation transition is

$$\langle t \rangle \sim \frac{\ell}{v} \sim \frac{\ell^2}{\mu \Delta E}. \quad (156)$$

On the other hand, diffusion over this time causes a spread in positions

$$\langle (\delta \ell)^2 \rangle \sim 2D \langle t \rangle = 2\mu k_B T \langle t \rangle, \quad (157)$$

where we make use of the Einstein relation $D = \mu k_B T$. Now (roughly speaking) since the molecule is moving in configuration space with typical velocity v , this spread in positions is equivalent to a variance in the time required to complete the transition to the de-activated state,

$$\langle (\delta t)^2 \rangle \sim \frac{\langle (\delta \ell)^2 \rangle}{v^2} \sim \frac{2\mu k_B T}{(\mu \Delta E / \ell)^2} \cdot \frac{\ell^2}{\mu \Delta E}. \quad (158)$$

If we express this as a fractional variance we find

$$\frac{\langle (\delta t)^2 \rangle}{\langle t \rangle^2} \sim \frac{2\mu k_B T}{(\mu \Delta E / \ell)^2} \cdot \frac{\ell^2}{\mu \Delta E} \cdot \left(\frac{\mu \Delta E}{\ell^2} \right)^2 \sim \frac{2k_B T}{\Delta E}. \quad (159)$$

Thus when we look at the variability of the lifetime in this model, the effective mobility μ and the magnitude ℓ of the structural change in the molecule drop out, and the reproducibility is just determined by the amount of energy that is dissipated in the de-activation transition. Indeed, comparing with the argument about multiple steps, our result here is the same as expected if the number of irreversible steps were $K \sim \Delta E / (2k_B T)$, consistent with the idea that each step must dissipate more than $k_B T$ in order to be effectively irreversible. To achieve a relative variance of 1/25 or 1/40 requires dropping $\sim 0.6 - 1$ eV (recall that $k_B T$ is 1/40 eV at room temperature), which is OK since the absorbed photon is roughly 2.5 eV.

Problem 29: Is there a theorem here? The above argument hints at something more general. Imagine that we have a molecule in some state, and we ask how long it takes to arrive at some other state. Assuming that the molecular dynamics is that of overdamped motion plus diffusion on some energy surface, can you show that the fractional variance in the time required for the motion is limited by the free energy difference between the two states?

How do we go about testing these different ideas? If saturation is important, one could try either by chemical manipulations or by genetic engineering to prolong the lifetime of rhodopsin and see if in fact the amplitude of the single photon response is buffered against these changes. If feedback is important, one could make a list of candidate feedback molecules and to manipulate the concentrations of these molecules. Finally, if there are multiple steps one could try to identify the molecular events associated with each step and perturb these events again either with chemical or genetic methods. All these are good ideas, and have been pursued by several groups.




FIG. 25 Variability in the single photon response with genetically engineered rhodopsins. (A) Wild type responses from mouse rods. Schematic shows the six phosphorylation sites, which are serine or threonine residues. In the remaining panels, we see responses when the number of phosphorylation sites has been reduced by mutating alanine, leaving five sites (B & C), two sites (D), one site (E), or none (F). From Doan et al (2007).

An interesting hint about the possibility of multiple steps in the rhodopsin shutoff is the presence of multiple phosphorylation sites on the opsin proteins. In mice, there are six phosphorylation sites, and one can genetically engineer organisms in which some or all of these sites are removed. At a qualitative level it's quite striking that even knocking out one of these sites produces a noticeable increase in the variability of the single photon responses, along with a slight prolongation of the mean response (Figs 25B & C). When all but one or two sites are removed, the responses last a *very* long time, and start to look like on/off switches with a highly variable time in the 'on' state (Figs 25D & E). When there are no phosphorylation sites, rhodopsin can still turn off, presumably as a result of binding another molecule (arrestin). But now the time to shutoff is broadly distributed, as one might expect if there were a single step controlling the




FIG. 26 Standard deviation in the integral of the single photon response, normalized by the mean. Results are shown as a function of the number of phosphorylation sites, from experiments as in Fig 25; error bars are standard errors of the mean. Solid line is $CV = 1/\sqrt{N_p + 1}$, where N_p is the number of phosphorylation sites. From Doan et al (2006).

Remarkably, if we examine the responses quantitatively, the variance of the single photon response seems to be inversely proportional the number of these sites, exactly as in the model where deactivation involved multiple steps, now identified with the multiple phosphorylations (Fig 26). This really is beautiful. One of the things that I think is interesting here is that, absent the discussion of precision and reproducibility, the multiple phosphorylation steps might just look like complexity for its own sake, the kind of thing that biologists point to when they want to tease physicists about our propensity to ignore details. In this case, however, the complexity seems to be the solution to a very specific physics problem.

Probably this section should end with some caveats. Do we really think the problem of reproducibility is solved?

A general review of the cGMP cascade in rods is given by Burns & Baylor (2001). Rieke & Baylor (1996) set out to understand the origins of the continuous noise in rods, but along the way provide a beautifully quantitative dissection of the enzymatic cascade; much of the discussion above follows theirs. For an explanation of how similarity to Rhodopsin (and other G-protein coupled receptors) drove the discovery of the olfactory receptors, see Buck (2004). For some general background on ion channels, you can try Aidley (see notes to Section 1.1), Johnston & Wu (1995), or Hille (2001). A starting point for learning about how different choices of channels shape the dynamics of responses in insect photoreceptors is the review by Weckström & Laughlin (1995). [There is much more to say here, and probably even some things left to do.]

Buck 2004: Unraveling the sense of smell. LB Buck, in *Les Prix Nobel: Nobel Prizes 2004*, T Frängsmyr, ed (Nobel Foundation, Stockholm, 2004).

Burns & Baylor 2001: Activation, deactivation and adaptation in vertebrate photoreceptor cells. ME Burns & DA Baylor, *Annu Rev Neurosci* **24**, 779–805 (2001).

Hille 2001: *Ion Channels of Excitable Membranes, 3rd Edition* B Hille (Sinauer, Sunderland MA, 2001).

Johnston & Wu 1995: *Foundations of Cellular Neurophysiology*. D Johnston & SM Wu (MIT Press Cambridge, 1995).

Rieke & Baylor 1996: Molecular origin of continuous dark noise in rod photoreceptors. F Rieke & DA Baylor, *Biophys J* **71**, 2553–2572 (1996).

Weckström & Laughlin 1995: Visual ecology and voltage gated ion channels in insect photoreceptors. M Weckström & SB Laughlin, *Trends Neurosci* **18**, 17–21 (1995).

Rieke & Baylor (1998a) provide a review of photon counting rods with many interesting observations, including an early outline of the problem of reproducibility. An early effort to analyze the signals and noise in enzymatic cascades is by Detwiler et al (2000). The idea that restricted, saturable domains can arise dynamically and tame the fluctuations in the output of the cascade is described by the same authors (Ramanathan et al 2005). For invertebrate photoreceptors, it seems that reproducibility of the response to single photons can be traced to positive feedback mechanisms that generate a stereotyped pulse of concentration changes, localized to substructures analogous to the disks in vertebrate rods (Pumir et al 2008).

Detwiler et al 2000: Engineering aspects of enzymatic signal transduction: Photoreceptors in the retina. PB Detwiler, S Ramanathan, A Sengupta & BI Shraiman, *Biophys J* **79**, 2801–2817 (2000).

Pumir et al 2008: A Pumir, J Graves, R Ranganathan & BI Shraiman, Systems analysis of the single photon response in invertebrate photoreceptors. *Proc Nat'l Acad Sci (USA)* **105**, 10354–10359 (2008).

Ramanathan et al 2005: G–protein–coupled enzyme cascades have intrinsic properties that improve signal localization and fidelity. S Ramanathan, PB Detwiler, AM Sengupta & BI Shraiman, *Biophys J* **88**, 3063–3071 (2005).

Rieke & Baylor 1998a: Single–photon detection by rod cells of the retina. F Rieke & DA Baylor, *Revs Mod Phys* **70**, 1027–1036 (1998).

One of the early, systematic efforts to test different models of reproducibility was by Rieke & Baylor (1998b). Many of the same ideas were revisited in mammalian rods by Field & Rieke (2002b), setting the stage for the experiments on genetic engineering of the phosphorylation sites by Doan et al (2006). More recent work from the same group explores the competition between the kinase and the arrestin molecule, which binds to the phosphorylated rhodopsin to terminate its activity, showing this competition influences both the mean and the variability of the single photon response (Doan et al 2009).

Doan et al 2007: Multiple phosphorylation sites confer reproducibility of the rod's single–photon responses. T Doan, A Mendez, PB Detwiler, J Chen & F Rieke, *Science* **313**, 530–533 (2006).

Doan et al 2009: Arrestin competition influences the kinetics and variability of the single–photon responses of mammalian rod photoreceptors. T Doan, AW Azevedo, JB Hurley & F Rieke, *J Neurosci* **29**, 11867–11879 (2009).

Field & Rieke 2002: Mechanisms regulating variability of the single photons responses of mammalian rod photoreceptors. GD Field & F Rieke, *Neuron* **35**, 733–747 (2002b).

Rieke & Baylor 1998a: Oring of reproducibility in the responses of retinal rods to single photons. F Rieke & DA Baylor, *Biophys J* **75**, 1836–1857 (1998).




FIG. 27 A schematic of the circuitry in the retina. **Fill in caption.**

of the integration is achieved in the very first step of processing, as many rod cells converge and form synapses onto onto a single bipolar cell, as shown schematically in Fig 27 [maybe also need a real retina?]

If each cell generates an output n_i that counts the number of photons that have arrived, then it's trivial that the total photon count is $n_{\text{total}} = \sum_i n_i$. The problem is that the cells don't generate integers corresponding to the number of photons counted, they generate currents which have continuous variations. In particular, we have seen that the *mean* current in response to a single photon has a peak of $I_1 \sim 1$ pA, but this rests on continuous background noise with an amplitude $\delta I_{\text{rms}} \sim 0.1$ pA. In a single cell, this means that the response to one photon stands well above the background, but if we try to sum the signals from many cells, we have a problem, as illustrated in Fig 28.

To make the problem precise, let's use x_i to denote the peak current generated by cell i . We have

$$x_i = I_1 n_i + \eta_i, \quad (160)$$

where n_i is the number of photons that are counted in cell i , and η_i is the background current noise; from what we have seen in the data, each η_i is chosen independently from a Gaussian distribution with a standard deviation δI_{rms} . If we sum the signals generated by all the cells, we obtain

$$x_{\text{total}} \equiv \sum_{i=1}^{N_{\text{cells}}} x_i = I_1 \sum_{i=1}^{N_{\text{cells}}} n_i + \sum_{i=1}^{N_{\text{cells}}} \eta_i \quad (161)$$

$$= I_1 n_{\text{total}} + \eta_{\text{eff}}, \quad (162)$$

where the effective noise is the sum of N_{cells} independent samples of the η_i , and hence has a standard deviation

$$\eta_{\text{eff}}^{\text{rms}} \equiv \sqrt{\langle \eta_{\text{eff}}^2 \rangle} = \sqrt{N_{\text{cells}}} \delta I_{\text{rms}}. \quad (163)$$

D. The first synapse, and beyond

This is a good moment to remember a key feature of the Hecht, Shlaer and Pirenne experiment, as described in Section I.A. In that experiment, observers saw flashes of light that delivered just a handful of photons spread over an area that includes many hundreds of photoreceptor cells. One consequence is that a single receptor cell has a very low probability of counting more than one photon, and this is how we know that these cells must respond to single photons. But, it must also be possible for the retina to add up the responses of these many cells so that the observer can reach a decision. Importantly, there is no way to know in advance which cells will get hit by photons, so if we (sliding ourselves into the positions of the observer's brain ...) want to integrate the multiple photon counts we have to integrate over all the receptors in the area covered by the flash. This integration might be the simplest computation we can imagine for a nervous system, just adding up a set of elementary signals, all given the same weight. In many retinas, a large part




FIG. 28 Simulation of the peak currents generated by $N = 500$ rod cells in response to a dim flash of light. At left, five of the cells actually detect a photon, each resulting in a current $I_1 \sim 1$ pA, while at right we see the response to a blank. All cells have an additive background noise, chosen from a Gaussian distribution with zero mean and standard deviation $\delta I_{\text{rms}} \sim 0.1$ pA. Although the single photon responses stand clearly above the background noise, if we simply add up the signals generated by all the cells, then at left we find a total current $I_{\text{total}} = 1.85$ pA, while at right we find $I_{\text{total}} = 3.23$ pA—the summed background noise completely overwhelms the signal.

The problem is that with $\delta I_{\text{rms}} \sim 0.1$ pA and $N_{\text{cells}} = 500$, we have $\eta_{\text{eff}}^{\text{rms}} \sim 2.24$ pA, which means that there is a sizable chance of confusing three or even five photons with a blank; in some species, the number of cells over which the system integrates is even larger, and the problem becomes even more serious. Indeed, in primates like us, a single ganglion cell (one stage after the bipolar cells; cf Fig 27) receives input from ~ 4000 rods, while on a very dark night we can see when just one in a thousand rods is captures a photon [should have refs for all this]. Simply put, summing the signals from many cells buries the clear single photon response under the noise generated by those cells which did not see anything. This can't be the right way to do things!

Before we start trying to do something formal, let's establish some intuition. Since the single photon signals are clearly detectable in individual rod cells, we could solve our problem by making a 'decision' for each cell—is there a photon present or not?—and then adding up the tokens that represent the outcome of our decision. Roughly speaking, this means passing each rod's signal through some fairly strong nonlinearity, perhaps so strong that it has as an output only a 1 or a 0, and then pooling these nonlinearly transformed signals. In contrast, a fairly standard schematic of what neurons are doing throughout the brain is adding up their inputs and then passing this sum through a nonlinearity (Fig 29).

So perhaps the problems of noise in photon counting are leading us to predict that this very first step of neural computation in the retina has to be different from this standard schematic. Let's try to do an honest calculation that makes this precise. [Is “nonlinearity” clear enough here?]

Formally, the problem faced by the system is as follows. We start with the set of currents generated by all the rod cells, $\{x_i\}$. We can't really be interested in the currents themselves. Ideally we want to know about what is happening in the outside world, but a first step would be to estimate the total number of photons that arrived, n_{total} . What is the best estimate we can make? To answer this, we need to say what we mean by “best.”

One simple idea, which is widely used, is that we want to make estimates which are as close as possible to the right answer, where closeness is measured by the mean square error. That is, we want to map the data $\{x_i\}$ into an estimate of n_{total} through some function $n_{\text{est}}(\{x_i\})$ such that

$$\mathcal{E} \equiv \left\langle [n_{\text{total}} - n_{\text{est}}(\{x_i\})]^2 \right\rangle \quad (164)$$

is as small as possible. To find the optimal choice of the function $n_{\text{est}}(\{x_i\})$ seems like a hard problem—maybe we have to choose some parameterization of this function, and then vary the parameters? In fact, we can solve this problem once and for all, which is part of the reason that this definition of ‘best’ is popular.

When we compute our average error, we are averaging over the joint distribution of the data $\{x_i\}$ and the actual




FIG. 29 Schematic of summation and nonlinearity in the initial processing of rod cell signals. At left, a conventional model in which many rods feed into one bipolar cell; the bipolar cell sums its inputs and passes the results through a saturating nonlinearity. At right, an alternative model, suggested by the problems of noise, in which nonlinearities precede summation.

photon count n_{total} . That is,

$$\begin{aligned} \mathcal{E} &\equiv \left\langle [n - n_{\text{est}}(\{x_i\})]^2 \right\rangle \\ &= \int \left[\prod_{i=1}^{N_{\text{cells}}} dx_i \right] \sum_{n_{\text{total}}} P(n, \{x_i\}) [n - n_{\text{est}}(\{x_i\})]^2 \end{aligned} \quad (165)$$

where, to simplify the notation, we drop the subscript total. Now to minimize the error we take the variation with respect to the function $n_{\text{est}}(\{x_i\})$ and set the result equal to zero. We have

$$\frac{\delta \mathcal{E}}{\delta n_{\text{est}}(\{x_i\})} = - \sum_n P(n, \{x_i\}) 2[n - n_{\text{est}}(\{x_i\})], \quad (166)$$

so setting this to zero gives (going through the steps carefully):

$$\sum_n P(n, \{x_i\}) n_{\text{est}}(\{x_i\}) = \sum_n P(n, \{x_i\}) n \quad (167)$$

$$n_{\text{est}}(\{x_i\}) \sum_n P(n, \{x_i\}) = \sum_n P(n, \{x_i\}) n \quad (168)$$

$$n_{\text{est}}(\{x_i\}) P(\{x_i\}) = \sum_n P(n, \{x_i\}) n \quad (169)$$

$$n_{\text{est}}(\{x_i\}) = \sum_n \frac{P(n, \{x_i\})}{P(\{x_i\})} n, \quad (170)$$

and, finally,

$$n_{\text{est}}(\{x_i\}) = \sum_n P(n|\{x_i\}) n. \quad (171)$$

Thus the optimal estimator is the mean value in the conditional distribution, $P(n|\{x_i\})$. Since we didn't use any special properties of the distributions, this must be true in general, as long as 'best' means to minimize mean square error. We'll use this result many times, and come back to the question of whether the choice of mean square error is a significant restriction.

Notice that the relevant conditional distribution is the distribution of photon counts given the rod cell currents. From a mechanistic point of view, we understand the opposite problem, that is, given the photon counts, we know how the currents are being generated. More precisely, we know that, given the number of photons in each cell, the currents will be drawn out of a probability distribution, since this is (implicitly) what we are saying when we write Eq (160). To make this explicit, we have

$$P(\{x_i\}|\{n_i\}) \propto \exp \left[-\frac{1}{2} \sum_{i=1}^{N_{\text{cells}}} \left(\frac{x_i - I_1 n_i}{\delta I_{\text{rms}}} \right)^2 \right]. \quad (172)$$

Again, this is a model that tells us how the photons generate currents. But the problem of the organism is to use the currents to draw inferences about the photons.

We expect that since the signals are noisy, this inference will be probabilistic, so really we would like to know $P(\{n_i\}|\{x_i\})$.

Problem 30: Just checking. Be sure that you understand the connection between Eq (172) and Eq (160). In particular, what assumptions are crucial in making the connection?

The problem of going from $P(\{x_i\}|\{n_i\})$ to $P(\{n_i\}|\{x_i\})$ is typical of the problems faced by organisms: given knowledge of how our sensory data is generated, how do we reach conclusions about what really is going on in the outside world? In a sense this is the same problem that we face in doing physics experiments. One could argue that what we have posed here is a very easy version of the real problem. In fact, we probably don't really care about the photon arrivals, but about the underlying light intensity, or more deeply about the identity and movements of the objects from which the light is being reflected. Still, this is a good start.

The key to solving these inference problems, both for organisms and for experimental physicists, is Bayes' rule. Imagine that we have two events A and B ; to be concrete, we could think of A as some data we observe, and B as a variable in the world that we really want to know. There is some probability $P(A, B)$ that both of these are true simultaneously, i.e. that we observe A and the world is in state B . In the usual view, the states of the world cause the data to be generated in our instruments, so we can say that the state of the world is chosen out of some distribution $P(B)$, and then given this state the data are generated, with some noise, and hence drawn out of the conditional distribution $P(A|B)$. By the usual rules of probability, we have

$$P(A, B) = P(A|B)P(B). \quad (173)$$

We could also imagine that we have just seen the data A , drawn out of some distribution $P(A)$, and then there must be some distribution of things happening in the world that are consistent with our observation. Formally,

$$P(A, B) = P(B|A)P(A). \quad (174)$$

But these are just two different ways of decomposing the joint distribution $P(A, B)$, and so they must be equal:

$$P(A, B) = P(B|A)P(A) = P(A|B)P(B) \quad (175)$$

$$P(B|A) = \frac{P(A|B)P(B)}{P(A)}. \quad (176)$$

This last equation is called Bayes' rule, and tells us what we need to know. It is useful to rewrite this, taking seriously the case where A refers to measurable data and B refers to the state of the world:

$$P(\text{world}|\text{data}) = \frac{P(\text{data}|\text{world})P(\text{world})}{P(\text{data})}. \quad (177)$$

Equation (177) is telling us that the probability of the world being in a certain state is proportional to the probability that this state could generate the data we have seen, but this is multiplied by the overall probability that the world can be in this state. This term often is referred to as the 'prior' probability, since it summarizes our knowledge prior to the observation of the data. Put another way, our inference about the world should be both consistent with the data we have observed in this one experiment *and* with any prior knowledge we might have from previous data.²³

Applied to our current problem, Bayes' rule tells us how to construct the probability distribution of photon counts given the rod currents:

$$P(\{n_i\}|\{x_i\}) = \frac{P(\{x_i\}|\{n_i\})P(\{n_i\})}{P(\{x_i\})}. \quad (178)$$

To make progress (and see how to use these ideas), let's start with the simple case of just one rod cell, so we can drop the indices:

$$P(n|x) = \frac{P(x|n)P(n)}{P(x)}. \quad (179)$$

To keep things really simple, let's just think about the case where the lights are very dim, so either there are zero photons or there is one photon, so that

$$P(1|x) = \frac{P(x|1)P(1)}{P(x)}, \quad (180)$$

and similarly for $P(0|x)$. In the denominator we have $P(x)$, which is the probability that we will see the current x , without any conditions on what is going on in the world. We get this by summing over all the possibilities,

$$P(x) = \sum_n P(x|n)P(n) \quad (181)$$

$$= P(x|1)P(1) + P(x|0)P(0), \quad (182)$$

where in the last step we use the approximation that the lights are very dim. Putting the terms together, we have

$$P(1|x) = \frac{P(x|1)P(1)}{P(x|1)P(1) + P(x|0)P(0)}. \quad (183)$$

Now we can substitute for $P(x|n)$ from Eq (172),

$$P(x|n) = \frac{1}{\sqrt{2\pi}(\delta I_{\text{rms}})^2} \exp\left[-\frac{(x - I_1 n)^2}{2(\delta I_{\text{rms}})^2}\right]. \quad (184)$$

Going through the steps, we have

$$P(1|x) = \frac{P(x|1)P(1)}{P(x|1)P(1) + P(x|0)P(0)} = \frac{1}{1 + P(x|0)P(0)/P(x|1)P(1)} \quad (185)$$

$$= \frac{1}{1 + [P(0)/P(1)] \exp\left[-\frac{(x)^2}{2(\delta I_{\text{rms}})^2} + \frac{(x - I_1)^2}{2(\delta I_{\text{rms}})^2}\right]} = \frac{1}{1 + \exp(\theta - \beta x)}, \quad (186)$$

where

$$\theta = \ln\left[\frac{P(0)}{P(1)}\right] + \frac{I_1^2}{2(\delta I_{\text{rms}})^2} \quad (187)$$

$$\beta = \frac{I_1}{(\delta I_{\text{rms}})^2}. \quad (188)$$

The result in Eq (186) has a familiar form—it is as if the two possibilities (0 and 1 photon) are two states of a physical system, and their probabilities are determined by a Boltzmann distribution; the energy difference between the two states shifts in proportion to the data x , and the temperature is related to the noise level in the system. In the present example, this analogy doesn't add much, essentially because the original problem is so simple, but we'll see richer cases later on in the course.

Equation (186) tells us that, if we observe a very small current x , the probability that there really was a photon present is small, $\sim e^{-\theta}$. As the observed current becomes larger, the probability that a photon was present goes up, and, gradually, as x becomes large, we become certain [$P(1|x) \rightarrow 1$]. To build the best estimator of n from this one cell, our general result tells us that we should compute the conditional mean:

$$n_{\text{est}}(x) = \sum_n P(n|x)n \quad (189)$$

$$= P(0|x) \cdot (0) + P(1|x) \cdot (1) \quad (190)$$

$$= P(1|x). \quad (191)$$

Thus, the Boltzmann-like result [Eq (186)] for the probability of a photon being counted is, in fact, our best

estimator of the photon count in this limit where photons are very rare. Further, in this limit one can show that the optimal estimator for the total photon count, which after all is the sum of the individual n_i , is just the sum of the individual estimators.

Problem 31: Summing after the nonlinearity. Show that the optimal estimator for the total number of photons is the sum of estimators for the photon counts in individual rods, provided that the lights are very dim and hence photons are rare. The phrasing here is deliberately vague—you should explore the formulation of the problem, and see exactly what approximations are needed to make things come out right.

The end result of our calculations is that the optimal estimator of photon counts really is in the form shown at the right in Fig 29: nonlinearities serve to separate signal from noise in each rod cell, and these ‘cleaned’ signals are summed. How does this prediction compare with experiment? Careful measurements in the mouse retina show that the bipolar cells respond nonlinearly even to very dim flashes of light, in the range where the rods see single photons and respond linearly, with two photons producing twice the response to one photon. The form of the nonlinearity is what we expect from the theory, a roughly sigmoidal function that suppresses noise and passes signals only above an amplitude threshold. Importantly, this nonlinearity is observed in one class of bipolar cells but not others, and this is the class that, on other grounds, one would expect is most relevant for processing of rod outputs at low light levels.

Looking more quantitatively at the experiments [show some of the data, perhaps replotted in different forms ... go back and look at the original papers and clean up this paragraph!], we can see discrete, single photon events in the bipolar cells. Although the details vary across organisms, in this retina, one bipolar cell collects input from ~ 20 rod cells, but the variance of the background noise is larger than in the lower vertebrates that we first saw in Fig 4. As a result, if we sum the rod inputs and pass them through the observed nonlinearity—as in the model at left in Fig 29—we would not be able to resolve the single photon events. Field and Rieke considered a family of models in which the nonlinearity has the observed shape but the midpoint (analogous to the threshold θ above) is allowed to vary, and computed the signal to noise ratio at the bipolar cell output for the detection of flashes corresponding to a mean count of $\sim 10^{-4}$ photons/rod cell, which is, approximately, the point at which we can barely see something on a moonless night. Changing the threshold by a factor of two changes the signal to noise ratio by factors of several hundred. The measured value

of the threshold is within 8% of the predicted optimal setting, certainly close enough to make us think that we are on the right track.

The discussion thus far has emphasized separating signals from noise by their amplitudes.²⁴ We also can see, by looking closely at the traces of current vs time, that signal and noise have different frequency content. This suggests that we could also improve the signal to noise ratio by filtering. It’s useful to think about a more general problem, in which we observe a time dependent signal $y(t)$ that is driven by some underlying variable $x(t)$; let’s assume that the response of y to x is linear, but noisy, so that

$$y(t) = \int_{-\infty}^{\infty} d\tau g(\tau)x(t - \tau) + \eta(t), \quad (192)$$

where $g(\tau)$ describes the response function and $\eta(t)$ is the noise. What we would like to do is to use our observations on $y(t)$ to estimate $x(t)$.

Problem 32: Harmonic oscillator revisited. Just to be sure you understand what is going in Eq (192), think again about the Brownian motion of a damped harmonic oscillator, as in Problem [***], but now with an external force $F(t)$,

$$m \frac{d^2x(t)}{dt^2} + \gamma \frac{dx(t)}{dt} + \kappa x(t) = F(t) + \delta F(t). \quad (193)$$

Show that

$$x(t) = \int_{-\infty}^{\infty} d\tau g(\tau)F(t - \tau) + \eta(t). \quad (194)$$

Derive an explicit expression for the Fourier transform of $g(\tau)$, and find $g(\tau)$ itself in the limit of either small or large damping γ .

Since the y is linearly related to x , we might guess that we can make estimates using some sort of linear operation. As we have seen already in the case of the rod currents, this might not be right, but let’s try anyway—we’ll need somewhat more powerful mathematical tools to sort out, in general, when linear vs nonlinear computations are the most useful. We don’t have any reason to prefer one moment of time over another, so we should do something that is both linear and invariant under time translations, which means that our estimate must be of the form

$$x_{\text{est}}(t) = \int_{-\infty}^{\infty} dt' f(t - t')y(t'), \quad (195)$$

²⁴ Need to be a little careful here, since the analysis from Fred’ lab actually involves applying the nonlinearity to voltages that have already been filtered. Presumably this will be clearer when I am pointing to the real data .. come back and fix this!

where $f(t)$ is the ‘filter’ that we hope will separate signal and noise. Following the spirit of the discussion above, we’ll ask that our estimate be as close as possible to the right answer in the sense of mean-square error. Thus, our task is to find the filter $f(t)$ that minimizes

$$\mathcal{E} = \left\langle \left[x(t) - \int_{-\infty}^{\infty} dt' f(t-t')y(t') \right]^2 \right\rangle. \quad (196)$$

In taking the expectation value of the mean-square error, we average over possible realizations of the noise and the variations in the input signal $x(t)$. In practice this averaging can also be thought of as including an average over time.²⁵ Thus we can also write

$$\mathcal{E} = \left\langle \int_{-\infty}^{\infty} dt \left[x(t) - \int_{-\infty}^{\infty} dt' f(t-t')y(t') \right]^2 \right\rangle. \quad (197)$$

This is useful because we can then pass to the Fourier domain. We recall that for any function $z(t)$,

$$\int_{-\infty}^{\infty} dt z^2(t) = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} |\tilde{z}(\omega)|^2, \quad (198)$$

and that the Fourier transform of a convolution is the product of transforms,

$$\int_{-\infty}^{\infty} dt e^{i\omega t} \int_{-\infty}^{\infty} dt' f(t-t')y(t') = \tilde{f}(\omega)\tilde{y}(\omega). \quad (199)$$

Putting things together, we can rewrite the mean-square error as

$$\mathcal{E} = \left\langle \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \left| \tilde{x}(\omega) - \tilde{f}(\omega)\tilde{y}(\omega) \right|^2 \right\rangle. \quad (200)$$

Now each frequency component of our filter $\tilde{f}(\omega)$ appears independently of all the others, so minimizing \mathcal{E} is straightforward. The result is that

$$\tilde{f}(\omega) = \frac{\langle \tilde{y}^*(\omega)\tilde{x}(\omega) \rangle}{\langle |\tilde{y}(\omega)|^2 \rangle}. \quad (201)$$

Problem 33: Details of the optimal filter. Fill in the steps leading to Eq (201). Be careful about the fact that $f(t)$ is real, and so the transform $\tilde{f}(\omega)$ is not arbitrary. Hint: think about positive and negative frequency components.

To finish our calculation, we go back to Eq (192), which in the frequency domain can be written as

$$\tilde{y}(\omega) = \tilde{g}(\omega)\tilde{x}(\omega) + \tilde{\eta}(\omega). \quad (202)$$

Thus

$$\langle \tilde{y}^*(\omega)\tilde{x}(\omega) \rangle = \tilde{g}^*(\omega)\langle |\tilde{x}(\omega)|^2 \rangle \quad (203)$$

$$\langle |\tilde{y}(\omega)|^2 \rangle = |\tilde{g}(\omega)|^2\langle |\tilde{x}(\omega)|^2 \rangle + \langle |\tilde{\eta}(\omega)|^2 \rangle. \quad (204)$$

If all of these variables have zero mean (which we can have be true just by choosing the origin correctly), then quantities such as $\langle |\tilde{x}(\omega)|^2 \rangle$ are the variances of Fourier components, which we know (see Appendix B) are proportional to power spectra. Finally, then, we can substitute into our expression for the optimal filter to find

$$\tilde{f}(\omega) = \frac{\tilde{g}^*(\omega)S_x(\omega)}{|\tilde{g}(\omega)|^2S_x(\omega) + S_\eta(\omega)}, \quad (205)$$

where, as before, S_x and S_η are the power spectra of x and η , respectively.

In the case that noise is small, we can let $S_\eta \rightarrow 0$ and we find

$$\tilde{f}(\omega) \rightarrow \frac{1}{\tilde{g}(\omega)}. \quad (206)$$

This means that, when noise can be neglected, the best way to estimate the underlying signal is just to invert the response function of our sensor, which makes sense. Notice that since \tilde{g} generally serves to smooth the time dependence of $y(t)$ relative to that of $x(t)$, the filter $\tilde{f}(\omega) \sim 1/\tilde{g}(\omega)$ undoes this smoothing. This is important because it reminds us that smoothing in and of itself does not set a limit to time resolution—it is only the combination of smoothing with noise that obscures rapid variations in the signal.

Guided by the limit of high signal to noise ratio, we can rewrite the optimal filter as

$$\tilde{f}(\omega) = \frac{1}{\tilde{g}(\omega)} \cdot \frac{|\tilde{g}(\omega)|^2S_x(\omega)}{|\tilde{g}(\omega)|^2S_x(\omega) + S_\eta(\omega)} \quad (207)$$

$$= \frac{1}{\tilde{g}(\omega)} \cdot \frac{SNR(\omega)}{1 + SNR(\omega)}, \quad (208)$$

where we identify the signal to noise ratio at each frequency, $SNR(\omega) = |\tilde{g}(\omega)|^2S_x(\omega)/S_\eta(\omega)$. Clearly, as the signal to noise ratio declines, so does the optimal filter—in the limit, if $SNR(\omega) = 0$, everything we find at frequency ω must be noise, and so it should be zeroed out if we want to minimize its corrupting effects on our estimates.

In the case of the retina, x is the light intensity, and y are the currents generated by the rod cells. When it’s very dark outside, the signal to noise ratio is low, so that

$$\tilde{f}(\omega) \rightarrow \frac{\tilde{g}^*(\omega)}{S_\eta(\omega)} \cdot S_x(\omega). \quad (209)$$

²⁵ More formally, if all the relevant random variations are ergodic, then averaging over the distributions and averaging over time will be the same.

The filter in this case has two pieces, one of which depends only on the properties of the rod cell,

$$\tilde{f}_1(\omega) = \frac{\tilde{g}^*(\omega)}{S_\eta(\omega)}, \quad (210)$$

and another piece that depends on the power spectrum of the time dependent light intensity, $S_x(\omega)$. With a bit more formalism we can show that this first filter, $\tilde{f}_1(\omega)$, has a universal meaning, so that if instead of estimating the light intensity itself, we try to estimate something else—e.g., the velocity of motion of an object across the visual field—then the first step in the estimation process is still to apply this filter. So, it is a natural hypothesis that this filter will be implemented near the first stages of visual processing, in the transfer of signals from the rods to the bipolar cells.




FIG. 30 Voltage responses of rod and bipolar cells in the salamander retina, compared with theory, from Rieke et al (1991). The theory is that the transmission from rod currents to bipolar cell voltage implements the optimal filter as in Eq (210). Measured responses are averages over many presentations of a flash at $t = 0$ that results in an average of five photons being counted. The predicted filter is computed from measured signal and noise properties of the rod cell, with no adjustable parameters.

Problem 34: Filtering the real rod currents. The raw data that were used to generate Fig 4 are available at <http://www.princeton.edu/~wbialek/PHY562/data.html>, in the file `rodcurrents.mat`. The data consist of 395 samples of the rod current in response to dim flashes of light. The data are sampled in 10 ms bins, and the flash is delivered in the 100th bin. If these ideas about filtering are sensible, we should be able to do a better job of discriminating between zero, one and two photons by using the right filter. Notice that filtering of a response that is locked to a particular moment in time is equivalent to taking a weighted linear combination of the currents at different times relative to the flash. Thus you can think of the current in response to one flash as a vector, and filtering amounts to taking the dot product of this vector with some template. As a first step, you should reproduce the results of Fig 4, which are based just on averaging points in the neighborhood of the peak. Under some conditions, the best

template would just be the average single photon response. How well does this work? What conditions would make this work best? Can you do better? These data are from experiments by FM Rieke and collaborators at the University of Washington, and thanks to Fred for making them available.

The idea that the rod/bipolar synapse implements an optimal filter is interesting not least because this leads us to a prediction for the dynamics of this synapse, Eq (210), which is written entirely in terms of the signal and noise characteristics of the rod cell itself. All of these properties are measurable, so there are no free parameters in this prediction.²⁶ To get some feeling for how these predictions work, remember that the noise in the rod cell has two components—the spontaneous isomerizations of rhodopsin, which have the same frequency content as the real signal, and the continuous background noise, which extends to higher frequency. If we have only the spontaneous isomerizations, then $S_\eta \sim |\tilde{g}|^2$, and we are again in the situation where the best estimate is obtained by ‘unsmoothing’ the response, essentially recovering sharp pulses at the precise moments when photons are absorbed. This unsmoothing, or high-pass filtering, is cut off by the presence of the continuous background noise, and the different effects combine to make \tilde{f}_1 a band-pass filter. By the time the theory was worked out, it was already known that something like band-pass filtering was happening at this synapse; among other things this speeds up the otherwise rather slow response of the rod. In Fig 30 we see a more detailed comparison of theory and experiment.

Problem 35: Optimal filters, more rigorously. Several things were left out of the optimal filter analysis above; let’s try to put them back here.

(a.) Assume that there is a signal $s(t)$, and we observe, in the simplest case, a noisy version of this signal, $y(t) = s(t) + \eta(t)$. Let the power spectrum of $s(t)$ be given by $S(\omega)$, and the power spectrum of the noise $\eta(t)$ be given by $N(\omega)$. Further, assume that both signal and noise have Gaussian statistics. Show that the distribution of signals given our observations is

$$P[s(t)|y(t)] = \frac{1}{Z} \exp \left[-\frac{1}{2} \int \frac{d\omega}{2\pi} \frac{|\tilde{s}(\omega) - \tilde{y}(\omega)|^2}{N(\omega)} - \frac{1}{2} \int \frac{d\omega}{2\pi} \frac{|\tilde{s}(\omega)|^2}{S(\omega)} \right]. \quad (211)$$

²⁶ We should be a bit careful here. The filter, as written, is not causal. Thus, to make a real prediction, we need to shift the filter so that it doesn’t have any support at negative times. To make a well defined prediction, we adopt the minimal delay that makes this work. One could perhaps do better, studying the optimal filtering problem with explicitly causal filters, and considering the tradeoff between errors and acceptable delays.

(b.) Show that the most likely function $\tilde{s}(\omega)$ given the data on y is also the best estimate in the least squares sense, and is given by

$$\tilde{s}_{\text{est}}^{(\text{nc})}(\omega) = \frac{S(\omega)}{S(\omega) + N(\omega)} \tilde{y}(\omega); \quad (212)$$

the superscript (nc) reminds us that this estimate does not respect causality. Show that this is consistent with Eq (205). Notice that you didn't assume the optimal estimator was linear, so you have shown that it is (!). Which of the assumptions here are essential in obtaining this result?

(c.) The non-causal estimator is Eq (212) is constructed by assuming that we have access to the entire function $y(t)$, with $-\infty < t < \infty$, as we try to estimate, for example $s(t = 0)$. If we want our estimator to be something that we can build, then we must impose causality: the estimate of $s(t)$ can be based only on the history $y_- \equiv y(t' < t)$. Another way of saying this is that we don't really know $y_+ \equiv y(t' > t)$, so we should average over this part of the trajectory. But the average should be computed in the distribution $P[y_+|y_-]$. To construct this, start by showing that

$$P[y_+, y_-] \equiv P[y(t)] = \frac{1}{Z_0} \exp \left[-\frac{1}{2} \int \frac{d\omega}{2\pi} \frac{|\tilde{y}(\omega)|^2}{S(\omega) + N(\omega)} \right]. \quad (213)$$

$$P[y_+, y_-] = \frac{1}{Z_0} \exp \left[-\frac{1}{2} \int_{-\infty}^0 dt \left| \int \frac{d\omega}{2\pi} e^{-i\omega t} \tilde{y}_-(\omega) \tilde{\psi}(\omega) \right|^2 - \frac{1}{2} \int_0^{\infty} dt \left| \int \frac{d\omega}{2\pi} e^{-i\omega t} (\tilde{y}_-(\omega) + \tilde{y}_+(\omega)) \tilde{\psi}(\omega) \right|^2 \right], \quad (217)$$

and that

$$P[y_+|y_-] \propto \exp \left[-\frac{1}{2} \int_0^{\infty} dt \left| \int \frac{d\omega}{2\pi} e^{-i\omega t} (\tilde{y}_-(\omega) + \tilde{y}_+(\omega)) \tilde{\psi}(\omega) \right|^2 \right]. \quad (218)$$

Explain why averaging over the distribution $P[y_+|y_-]$ is equivalent to imposing the "equation of motion"

$$\int \frac{d\omega}{2\pi} e^{-i\omega t} (\tilde{y}_-(\omega) + \tilde{y}_+(\omega)) \tilde{\psi}(\omega) = 0 \quad (219)$$

at times $t > 0$.

(f.) Write the non-causal estimate Eq (212) in the time domain as

$$s_{\text{est}}^{(\text{nc})}(t) = \int \frac{d\omega}{2\pi} e^{-i\omega t} \tilde{\psi}^*(\omega) \tilde{\psi}(\omega) \tilde{y}(\omega). \quad (220)$$

But the combination $\tilde{\psi}(\omega) \tilde{y}(\omega)$ is the Fourier transform of $z(t)$, which is the convolution of $\psi(t)$ with $y(t)$. Show that Eq (219) implies that the average of $z(t)$ is the distribution $P[y_+|y_-]$ vanishes for $t > 0$, and hence the averaging over y_+ is equivalent to replacing

$$\tilde{\psi}(\omega) \tilde{y}(\omega) \rightarrow \int_{-\infty}^0 d\tau e^{i\omega\tau} \int \frac{d\omega'}{2\pi} \tilde{\psi}(\omega') \tilde{y}(\omega') e^{-i\omega'\tau} \quad (221)$$

in Eq (212). Put all the pieces together to show that there is a causal estimate of $s(t)$ which can be written as

$$s_{\text{est}}(t) = \int \frac{d\omega}{2\pi} e^{-i\omega t} \tilde{k}(\omega) \tilde{y}(\omega), \quad (222)$$

where

$$\tilde{k}(\omega) = \tilde{\psi}(\omega) \int_0^{\infty} d\tau e^{i\omega\tau} \int \frac{d\omega'}{2\pi} e^{-i\omega'\tau} S(\omega') \tilde{\psi}^*(\omega'). \quad (223)$$

Verify that this filter is causal.

(d.) Recall that when we discuss causality, it is useful to think about the frequency ω as a complex variable. Explain why we can write

$$\frac{1}{S(\omega) + N(\omega)} = |\tilde{\psi}(\omega)|^2, \quad (214)$$

where $\tilde{\psi}(\omega)$ has no poles in the upper half of the complex ω plane. Verify that, with this decomposition,

$$\psi(t) = \int \frac{d\omega}{2\pi} e^{-i\omega t} \tilde{\psi}(\omega) \quad (215)$$

is causal, that is $\psi(t < 0) = 0$. Consider the case where the signal has a correlation time τ_c , so that $S(\omega) = 2\sigma^2\tau_c/[1 + (\omega\tau_c)^2]$, and the noise is white $N(\omega) = N_0$; construct $\tilde{\psi}(\omega)$ explicitly in this case.

(e.) Putting Eq (213) together with Eq (214), we can write

$$P[y_+, y_-] = \frac{1}{Z_0} \exp \left[-\frac{1}{2} \int \frac{d\omega}{2\pi} \left| \tilde{y}(\omega) \tilde{\psi}(\omega) \right|^2 \right]. \quad (216)$$

Show that

filtering. In the other, the same separation is achieved by a static nonlinearity, applied in practice to a linearly filtered signal. Presumably there is some more general nonlinear dynamic transformation that really does the best job. We expect that the proper mix depends on the detailed spectral structure of the signals and noise, and on the relative amplitudes of the signal and noise, which might be why the different effects are clearest in retinas from very different species. Indeed, there is yet another approach which emphasizes that the dynamic range of neural outputs is limited, and that this constrains how much information the second order neuron can provide about visual inputs; filters and nonlinearities can be chosen to optimize this information transmission across a wide range of background light intensities, rather than focusing only on the detectability of the dimmest lights. This approach has received the most attention in invertebrate retinas, such as the fly that we met near the end of Section I.A, and we will return to these ideas in Chapter 4. It would be nice to see this all put together correctly, and this is an open problem, surely with room for some surprises.

So far we have followed the single photon signal from the single rhodopsin molecule to the biochemical network that amplifies this molecular event into a macroscopic current, and then traced the processing of this electrical signal as it crosses the first synapse. To claim that we have said anything about *vision*, we have to at least follow the signal out of the retina and on its way to the brain. [By now we should have said more about retinal anatomy—optic nerve, made up of the axons from 'retinal ganglion cells,' and the stereotyped action potentials

It is worth noting that we have given two very different analyses. In one, signals and noise are separated by linear

that propagate along these axons. Should also discuss techniques for picking up the signals, up to current work with electrode arrays. Show a modern figure, e.g. from Berry's lab.]

The classic experiments on single photon responses in retinal ganglion cells were done well before it was possible to measure the responses of single rods. The spikes from single ganglion cells are relatively easy to record, and one can try to do something like the Hecht, Shlaer and Pirenne experiment, but instead of “seeing” (as in Fig 2), you just ask if you can detect the spikes. There were a number of hints in the data that a single absorbed photon generated more than one spike, so some care is required. As shown in Fig 31, there are neurons that seem to count by threes—if you wait for three spikes, the probability of seeing is what you expect for setting a threshold of $K = 1$ photon, if you wait for six spikes it is as if $K = 2$, and so on. This simple linear relation between photons and spikes also makes it easy to estimate the rate of spontaneous photon-like events in the dark. Note that if photons arrive as a Poisson process, and each photon generates multiple spikes, then the spikes are *not* a Poisson process; this idea of Poisson events driving a second point process to generate non-Poisson variability has received renewed attention in the context of gene expression, where the a single messenger RNA molecule (perhaps generated from a Poisson process) can be translated to yield multiple protein molecules.




FIG. 31 A frequency of seeing experiment with spikes, from Barlow et al (1971). Recording from a single retinal ganglion cell, you can say you “saw” a flash when you detect 3, 6, 9, ... or more spikes within a small window of time (here, 200 ms). The probability of reaching this criterion is plotted vs the log of the flash intensity, as in the original Hecht, Shlaer and Pirenne experiments (Fig 2), but here the intensity is adjusted to include a background rate of photon-like events (“dark light”). Curves are from Eq (2), with the indicated values of the threshold K . Notice that three spikes corresponds to one photon.

Problem 36: Poisson-driven bursts. A characteristic feature of events drawn out of a Poisson process is that if we count the number of events, the variance of this number is equal to the mean. Suppose that each photon triggers exactly b spikes. What is the ratio of variance to mean (sometimes called the Fano factor) for spike counts in response to light flashes of fixed intensity? Suppose the the burst of spikes itself is a Poisson process, with mean b . Now what happens to the variance/mean ratio?

Before tracing the connections between individual spikes and photons, it was possible to do a different experiment, just counting spikes in response to flashes of different intensities, and asking what is the smallest value of the difference ΔI such that intensities I and $I + \Delta I$ can be distinguished reliably. The answer, of course, depends on the background intensity I [show figure from Barlow (1965)?]. For sufficiently small I , the just noticeable different ΔI is constant. For large I , one finds $\Delta I \propto I$, so the just noticeable fractional change in intensity is constant; this is common to many perceptual modalities, and is called Weber’s law. At intermediate intensities one can see $\Delta I \propto \sqrt{I}$. This last result, predicted by Rose and de Vries (cf Section 1.1), is what you expect if detecting a change in intensity just requires discriminating against the Poisson fluctuations in the arrival of photons. At high intensities, we are counting many photons, and probably the system just can’t keep up; then fluctuations in the gain of the response dominate, and this can result in Weber’s law. At the lowest intensities, the photons delivered by the flash are few in comparison with the thermal isomerizations of Rhodopsin, and this constant noise source sets the threshold for discrimination. Happily, the rate of spontaneous isomerizations estimated from these sorts of experiments agrees with other estimates, including the (much later) direct measurements on rod cells discussed previously. This work on discrimination with neurons also is important because it represents one of the first efforts to connect the perceptual abilities of whole organisms with the response of individual neurons.

If retinal ganglion cells generate three spikes for every photon, lights wouldn’t need to be very bright before the cells should be generating thousands of spikes per second, and this is impossible—the spikes themselves are roughly one millisecond in duration, and all neurons have a ‘refractory period’ that defines a minimum time (like a hard core repulsion) between successive action potentials. The answer is something we have seen already in the voltage responses of fly photoreceptors (Fig 13): as the background light intensity increases, the retina adapts and turns down the gain, in this case generating fewer spikes per photon. Of course this takes some time, so if we

suddenly expose the retina to a bright light there is very rapid spiking, which then adapts away to a much slower rate. [Need a figure about light/dark adaptation.] If we imagine that our perceptions are driven fairly directly by the spikes, then our impression of the brightness of the light should similarly fade away. This is true not just for light (as you experience whenever you walk outside on a bright sunny day); almost all constant sensory inputs get adapted away—think about the fact that you don't feel the pressure generated by your shoes a few minutes after you tie them. But there are more subtle issues as well, involving the possibility that the coding strategy used by the retina adapts to the whole distribution of inputs rather than just the mean; this is observed, and many subsequent experiments are aimed at understanding the molecular and cellular mechanisms of these effects. The possibility that adaptation serves to optimize the efficiency of coding continuous signals into discrete spikes is something we will return to in Chapter 4.

The problem of photon counting—or any simple detection task—also hides a deeper question: how does the brain “know” what it needs to do in any given task? Even in our simple example of setting a threshold to maximize the probability of a correct answer, the optimal observer must at least implicitly acquire knowledge of the relevant probability distributions. Along these lines, there is more to the ‘toad cooling’ experiment than a test of photon counting and dark noise. The retina has adaptive mechanisms that allow the response to speed up at higher levels of background light, in effect integrating for shorter times when we can be sure that the signal to noise ratio will be high. The flip side of this mechanism is that the retinal response slows down dramatically in the dark [connect back to photoreceptor responses; a figure here would be good, including τ vs I relevant to Aho et al]. In moderate darkness (dusk or bright moonlight) the slowing of the retinal response is reflected directly in a slowing of the animal's behavior. It is as if the toad experiences an illusion because images of its target are delayed, and it strikes at the delayed image. It is worth emphasizing that we see a closely related illusion.

Problem 37: Knowing where to look. Give a problem to illustrate the role of uncertainty in reducing performance.

Imagine watching a pendulum swinging while wearing glasses that have a neutral density filter over one eye, so the mean light intensity in the two eyes is different. The dimmer light results in a slower retina, so the signals from the two eyes are not synchronous, and recall that differences in the images between our right and left eyes

are cues to the depth of an object. As we try to interpret these signals in terms of motion, we find that even if the pendulum is swinging in a plane parallel to the line between our eyes, what we see is motion in 3D. The magnitude of the apparent depth of oscillation is related to the neutral density and hence to the slowing of signals in the ‘darker’ retina. This is called the Pulfrich effect.

If the pattern of delay vs light intensity continued down to the light levels in the darkest night, it would be a disaster, since the delay would mean that the toad inevitably strikes behind the target! In fact, the toad does not strike at all in the first few trials of the experiment in dim light, and then strikes well within the target. It is hard to escape the conclusion that the animal is learning about the typical velocity of the target and then using this knowledge to extrapolate and thereby correct for retinal delays.²⁷ Thus, performance in the limit where we count photons involves not only efficient processing of these small signals but also learning as much as possible about the world so that these small signals become interpretable.

If you'd like a general overview of the retina, a good source is Dowling (1987). For the experiments on nonlinear summation at the rod–bipolar synapse, along with a discussion of the theoretical issues of noise and reliability, see Field & Rieke (2002a). The analysis of optimal filtering is presented in Bialek & Owen (1990) and Rieke et al (1991). For a discussion how our experience of a dark night translates into photons per rod per second, see Walraven et al (1990).

- Bialek & Owen 1990:** Temporal filtering in retinal bipolar cells: Elements of an optimal computation? W Bialek & WG Owen, *Biophys J* **58**, 1227–1233 (1990).
- Dowling 1987:** *The Retina: An Approachable Part of the Brain* JE Dowling (Harvard University Press, Cambridge, 1987).
- Field & Rieke 2002a:** Nonlinear signal transfer from mouse rods to bipolar cells and implications for visual sensitivity. GD Field & F Rieke, *Neuron* **34**, 773–785 (2002).
- Rieke et al 1991:** Optimal filtering in the salamander retina. F Rieke, WG Owen & W Bialek, in *Advances in Neural Information Processing 3*, R Lippman, J Moody & D Touretzky, eds, pp 377–383 (Morgan Kaufmann, San Mateo CA, 1991).
- Walraven et al 1990:** The control of visual sensitivity. J Walraven, C Enroth–Cugell, DC Hood, DIA MacLeod & JL Schnapf, in *Visual Perception: The Neurophysiological Foundations*, L Spillmann & SJ Werner, eds, pp 53–101 (Academic Press, San Diego, 1990).

The classic presentations of filtering, estimation and prediction are by Kolmogorov (1939, 1941) and Wiener (1949). The long problem about optimal filtering is based on Potters & Bialek (1994).

Kolmogoroff 1939: Sur l'interpolation et extrapolations des suites stationnaires. A Kolmogoroff, *C R Acad Sci Paris* **208**, 2043–2045 (1939).

²⁷ As far as I know there are no further experiments that probe this learning more directly, e.g. by having the target move at variable velocities.

Kolmogorov 1941: Interpolation and extrapolation of stationary random sequences (in Russian). AN Kolmogorov, *Izv Akad Nauk USSR Ser Mat* **5**, 3–14 (1941). English translation in *Selected Works of AN Kolmogorov, Vol II*, AN Shiryagev, ed, pp 272–280 (Kluwer Academic, Dordrecht, 1992).

Potters & Bialek 1994: Statistical mechanics and visual signal processing. M Potters & W Bialek, *J Phys I France* **4**, 1755–1775 (1994); arXiv:cond-mat/9401072 (1994).

Wiener 1949: *Extrapolation, Interpolation and Smoothing of Time Series* N Wiener (Wiley, New York, 1949).

The idea of maximizing information transmission across the first visual synapse is something we will discuss at greater length in Chapter 4. Still, you might like to look ahead, so here are some references to how these ideas developed in the context of fly vision.

Hateren 1992: Real and optimal neural images in early vision. JH van Hateren, *Nature* **360**, 68–70 (1992).

Laughlin 1981: A simple coding procedure enhances a neuron’s information capacity. SB Laughlin, *Z Naturforsch* **36c**, 910–912 (1981).

Srinivasan et al 1982: Predictive coding: A fresh view of inhibition in the retina. MV Srinivasan, SB Laughlin & A Dubs, *Proc R Soc Lond Ser B* **216**, 427–459 (1982).

The classic paper about single photon responses in retinal ganglion cells is Barlow et al (1971); it has quite a lot of detail, and still makes great reading. [Mastronade 1983?; might also need pointers to more modern recordings] The idea that single molecular events can drive bursts, generating non-Poisson statistics, reappears thirty years later in the context of gene expression; see for example Ozbudak et al (2002). The early papers on intensity discrimination using spikes from single neurons are Barlow (1965) and Barlow & Levick (1969); see also the even earlier work from FitzHugh (1957, 1958).

Barlow 1965: Optic nerve impulses and Weber’s law. HB Barlow, *Cold Spring Harb Symp Quant Biol* **30**, 539–546 (1965).

Barlow & Levick 1969: Three factors limiting the reliable detection of light by retinal ganglion cells of the cat. HB Barlow & WR Levick, *J Physiol (Lond)* **200**, 1–24 (1969).

Barlow et al 1971: Responses to single quanta of light in retinal ganglion cells of the cat. HB Barlow, WR Levick & M Yoon, *Vision Res Suppl* **3**, 87–101 (1971).

FitzHugh 1957: The statistical detection of threshold signals in the retina. R FitzHugh, *J Gen Physiol* **40**, 925–948 (1957).

FitzHugh 1958: A statistical analyzer for optic nerve messages. R FitzHugh, *J Gen Physiol* **41**, 675–692 (1958).

Ozbudak et al 2002: Regulation of noise in the expression of a single gene. E Ozbudak, M Thattai, I Kurtser, AD Grossman & A van Oudenaarden, *Nature Gen* **31**, 69–73 (2002).

The observation that neurons gradually diminish their response to constant stimuli goes back to Adrian’s first experiments recording the spikes from single cells; he immediately saw the connection to the fading of our perceptions when inputs are constant, and this sort of direct mapping from neural responses to human experience is now the common language we use in thinking about the brain and mind. An early paper about adaptation to the distribution of inputs is Smirnakis et al (1997). Since then a number of papers have explored more complex versions of this adaptation, as well as trying to tease apart the underlying mechanisms; some examples are Rieke (2001), Kim & Rieke (2001, 2003), and Baccus & Meister (2002).

Adrian 1928: *The Basis of Sensation* ED Adrian (Christopher’s, London, 1928).

Baccus & Meister 2002: Fast and slow adaptation in retinal circuitry. SA Baccus & M Meister, *Neuron* **36**, 909–919 (2002).

Kim & Rieke 2001: Temporal contrast adaptation in the input and output signals of salamander retinal ganglion cells. KJ Kim & F Rieke, *J Neurosci* **21**, 287–299 (2001).

Kim & Rieke 2003: Slow Na⁺ inactivation and variance adaptation in salamander retinal ganglion cells. *J Neurosci* **23**, 1506–1515 (2003).

Rieke 2001: Temporal contrast adaptation in salamander bipolar cells. F Rieke, *J Neurosci* **21**, 9445–9454 (2001).

Smirnakis et al 1997: Adaptation of retinal processing to image contrast and spatial scale. S Smirnakis, MJ Berry II, DK Warland, W Bialek & M Meister, *Nature* **386**, 69–73 (1997).

There is a decent demonstration of the Pulfrich effect available on the web (Newbold 1999). The experiments on reaction times in toads and the connection to retinal delays are from the work of Aho et al (1993).

Aho et al 1993: Visual performance of the toad (*Bufo bufo*) at low light levels: Retinal ganglion cell responses and prey-catching accuracy. A-C Aho, K Donner, S Helenius, LO Larsen & T Reuter, *J Comp Physiol A* **172**, 671–682 (1993).

Newbold 1999: The Pulfrich illusion. M Newbold, <http://dogfeathers.com/java/pulfrich.html> (1999).

E. Perspectives

What have we learned from all of this? I think the first thing to notice is that we have at least one example of a real biological system that is susceptible to the sorts of reproducible, quantitative experiments that we are used to in the rest of physics. This is not obvious, and runs counter to some fairly widespread prejudices. Although things can get complicated,²⁸ it does seem that, with care, we can speak precisely about the properties of cells in the retina, not just on average over many cells but cell by cell, in enough detail that even the noise in the cellular response itself is reproducible from cell to cell, organism to organism. It’s important that all of this is not guaranteed—removing cells from their natural milieu can be traumatic, and every trauma is different. If you dig into the original papers, you’ll see glimpses of the many things which experimentalists need to get right in order to achieve the level of precision that we have emphasized in our discussion—the things one needs to do in order to turn the exploration of living systems into a physics experiment.

The second point is that the performance of these biological systems—something which results from mechanisms of incredible complexity—really is determined by the physics of the “problem” that the system has been selected to solve. If you plan on going out in the dark of night, there is an obvious benefit to being able to detect dimmer sources of light, and to making reliable discriminations among subtly different intensities and, ultimately, different spatiotemporal patterns. You can’t do

²⁸ We have not explored, for example, the fact that there are many kinds of ganglion cells.

better than to count every photon, and the reliability of photon counting by the system as a whole can't be better than the limits set by noise in the detector elements. The fact that real visual systems reach these limits is extraordinary.

The last point concerns the nature of the explanations that we are looking for. We have discussed the currents generated in response to single photons, the filter characteristics and nonlinearities of synapses, and the spiking outputs of ganglion cells, and in each case we can ask why these properties of the system are as we observe them to be. Importantly, we can ask analogous questions about a wide variety of systems, from individual molecules to the regulation of gene expression in single cells to the dynamics of neural networks in our brains. What are we doing when we look for an "explanation" of the data?

When we ask "why" in relation to a biological system, we can imagine (at least) two very different kinds of answers.²⁹ First, we could plunge into the microscopic mechanisms. As we have seen in looking at the dynamics of biochemical amplification in the rod cell, what we observe as functional behavior of the system as a whole depends on a large number of parameters: the rates of various chemical reactions, the concentrations of various proteins, the density of ion channels in the membrane, the binding energies of cGMP to the channel, and so on. To emphasize the obvious, these are not fundamental constants. In a very real sense, almost all of these

parameters are under the organism's control.

Our genome encodes hundreds of different ion channels, and the parameters of the rod cell would change if it chose to read out the instructions for making one channel rather than another. Further, the cell can make more or less of these proteins, again adjusting the parameters of the system essentially by changing the concentrations of relevant molecules. A similar line of argument applies to other components of the system (and many other systems), since many key molecules are members of families with slightly different properties, and cells choose which member of the family will be expressed. More subtly, many of these molecules can be modified, e.g. by covalent attachment of phosphate groups as with the shut-off of rhodopsin, and these modifications provide another pathway for adjusting parameters. Thus, saying that (for example) the response properties of the rod cell are determined by the parameters of a biochemical network is very different from saying that the absorption spectrum of hydrogen is determined by the charge and mass of the electron—we would have to go into some alternative universe to change the properties of the electron, but most of the parameters of the biochemical network are under the control of the cell, and these parameters can and do change in response to other signals.

An explanation of functional behavior in microscopic terms, then, may be correct but somehow unsatisfying. Further, there may be more microscopic parameters than phenomenological parameters, and this may be critical in allowing the system to achieve nearly identical functional behaviors via several different mechanisms. But all of this casts doubt on the idea that we are 'explaining' the functional behavior in molecular terms.

A second, very different kind of explanation is suggested by our discussion of the first synapse in vision, between the rod and bipolar cells. In that discussion (Section I.D), we promoted the evidence of near optimal performance at the problem of photon counting into a principle from which the functional properties of the system could be derived. In this view, the system is the way it is because evolution has selected the best solution to a problem that is essential in the life of the organism. This principle doesn't tell us how the optimum is reached, but it can predict the observable behavior of the system. Evidently there are many objections to this approach, but of course it is familiar, since many different ideas in theoretical physics can be formulated as variational principles, from least action in classical mechanics to the minimization of free energy in equilibrium thermodynamics, among others.

Organizing our thinking about biological systems around optimization principles tends to evoke philosophical discussions, in the pejorative sense that scientists use this term. I would like to avoid discussions of this flavor. If we are going to suggest that "biological systems maximize X" is a principle, then rather than having ev-

²⁹ My colleague Rob de Ruyter van Steveninck has an excellent way of talking about closely related issues. He once began a lecture by contrasting two different questions: Why is the sky blue? Why are trees green?. The answer to the first question is a standard part of a good, high level course on electromagnetism: when light scatters from a small particle—and molecules in the atmosphere are much smaller than the wavelength of light—the scattering is stronger at shorter wavelengths; this is called Rayleigh scattering. Thus, red light (long wavelengths) moves along a more nearly straight path than does blue light (short wavelength). The light that we see, which has been scattered, therefore has been enriched in the blue part of the spectrum, and this effect is stronger if we look further away from the sun. So, the sky is blue because of the way in which light scatters from molecules. We can answer the question about the color of trees in much the same way that we answered the question about the color of the sky: leaves contain a molecule called chlorophyll, which is quite a large molecule compared with the oxygen and nitrogen in the air, and this molecule actually absorbs visible light; the absorption is strongest for red and blue light, so what is reflected back to us is the (intermediate wavelength) green light. Unlike the color of the sky, the color of trees could have a different explanation. Imagine trees of other colors—blue, red, perhaps even striped. Microscopically, this could happen because their leaves contain molecules other than chlorophyll, or even molecules related to chlorophyll but with slightly different structures. But trees of different colors will compete for resources, and some will grow faster than others. The forces of natural selection plausibly will cause one color of tree to win out over the others. In this sense, we can say that trees are green because green trees are more successful, or more fit in their environment.

everyone express their opinion about whether this is a good idea, we should discipline ourselves and insist on criteria that allow such claims to be meaningful and predictive. First, we have to understand why X can't be arbitrarily large—we need to have a theory which defines the physical limits to performance. Second, we should actually be able to measure X, and compare its value with this theoretical maximum. Finally, maximizing X should generate some definite predictions about the structure and dynamics of the system, predictions that can be tested in independent, quantitative experiments. In what follows, we'll look at three different broad ideas about what X might be, and hopefully we'll be able to maintain the discipline that I have outlined here. Perhaps the most important lesson from the example of photon counting is that we can carry through this program and maintain contact with real data. The challenge is to choose principles (candidate Xs) that are more generally applicable than the very specific idea that the retina maximizes the reliability of seeing on a dark night.

II. NOISE ISN'T NEGLIGIBLE

The great poetic images of classical physics are those of determinism and clockwork. In a clock, not only the output but also the internal mechanisms are models of precision. Strikingly, life seems very different. Interactions between molecules involve energies of just a few times the thermal energy. Biological motors, including the molecular components of our muscles, move in elementary steps that are on the nanometer scale, driven forward by energies that are larger than the thermal energies of Brownian motion, but not much larger. Crucial signals inside cells often are carried by just a handful of molecules, and these molecules inevitably arrive randomly at their targets. Human perception can be limited by noise in the detector elements of our sensory systems, and individual elements in the brain, such as the synapses that pass signals from one neuron to the next, are surprisingly noisy. How do the obviously reliable functions of life emerge from under this cloud of noise? Are there principles at work that select, out of all possible mechanisms, the ones that maximize reliability and precision in the presence of noise?

In this Chapter, we will take a tour of various problems involving noise in biological systems. I should admit up front that this is a topic that always has fascinated me, and I firmly believe that there is something deep to be found in exploration of these issues. We will see the problems of noise in systems ranging from the behavior of individual molecules to our subjective, conscious experience of the world. In order to address these questions, we will need a fair bit of mathematical apparatus, rooted in the ideas of statistical physics. I hope that, armed with this apparatus, you will have a deeper view of many beautiful phenomena, and a deeper appreciation for the problems that organisms have to solve.

A. Molecular fluctuations and chemical reactions

In order to survive, living organisms must control the rates of many chemical reactions. Fundamentally, all reactions happen because of fluctuations. More strongly, chemical reactions are a non-perturbative consequence of molecular fluctuations. You all learned, perhaps even in high school, that the rates of chemical reactions obey the Arrhenius law, $k \propto e^{-E_{\text{act}}/k_B T}$, where E_{act} is the activation energy. We also know that $k_B T$ measures the mean square amplitude of fluctuations, for example in the velocities of atoms. Thus, chemical reaction rates are $\sim e^{-1/g}$, where g is the strength of the fluctuations. If we start by imagining a world in which there are no fluctuations, we can add them in piece by piece, but there is no way to get a chemical reaction rate as a perturbative series in g . Chemical reactions are so commonplace that we sometimes forget just how nontrivial they are from a




FIG. 32 The simplest model of a chemical reaction. Along some molecular coordinate x , the potential energy $V(x)$ has two minima separated by a barrier. The height of the barrier is the “activation energy” E_{act} , which we expect will determine the rate of the reaction through the Arrhenius law, $k \propto e^{-E_{\text{act}}/k_B T}$.

theoretical point of view. Indeed, as I verify every year, few of the students in my course have ever seen an honest calculation that gives the Arrhenius law as a result, although they have all heard vague arguments about the Boltzmann probability of being on top of the barrier. So, our first order of business is to see how the Arrhenius law emerges, as an asymptotic result, for some real dynamical model. Only once we have this more solid understanding will we be ready to look at what might be special regarding the control of chemical reaction rates in biological systems.

Let us consider the simplest case, shown in Fig 32. Here the molecules of interest are described by a single coordinate x , and the potential energy $V(x)$ as a function of this coordinate has two wells that we can identify as reactant and product structures. Let’s assume that motions along this coordinate are overdamped, so inertia is negligible.³⁰ Since the molecule is surrounded by an environment at temperature T , we really want to describe Brownian motion in this potential. So, the equation of motion is³¹

$$\gamma \frac{dx}{dt} = -\frac{dV(x)}{dx} + \zeta(t), \quad (224)$$

³⁰ This really is just a simplifying assumption. We can also do everything in the case where inertia is significant, and none of the main results will be different. More precisely, we are going to go far enough to show that the Arrhenius law $k = A e^{-E_{\text{act}}/k_B T}$ is true, and that the activation energy E_{act} corresponds to our intuition. The neglect of inertia would only change the prefactor A , which is in any case much more difficult to calculate.

³¹ For background on the description of random functions of time, see Appendix A.2.

where γ is the friction or drag coefficient, and the random or Langevin force $\zeta(t)$ reflects the random influences of all the other degrees of freedom in the system; to insure that the system eventually comes to equilibrium at temperature T we must have

$$\langle \zeta(t)\zeta(t') \rangle = 2\gamma k_B T \delta(t - t'). \quad (225)$$

The challenge is to see if we can extract from these dynamics some approximate result which corresponds to our intuition about chemical reactions, and in particular gives us the exponential dependence of the rate on the temperature.

[Perhaps should add some discussion of the “reaction coordinate” concept. On the other hand, one could say that we are just doing the simplest case, which is one dimensional, in which case there is no need for apologies, just generalization later. Advice welcome.]

When we solve Eq (224), what we get is the coordinate as a function of time. What features of this trajectory correspond to the reaction rate k ? If there really are only two states in the sense of chemical kinetics, then trajectories should look like those in Fig 33. Specifically, we should see that trajectories spend most of their time in one potential well or the other, punctuated by rapid jumps between the wells. More precisely, there should be a clear separation of time scales between the dynamics within each well and the typical time between jumps. Further, if we look at the times spent in each well, between jumps, these times should be drawn from an exponential distribution, $P(t) = ke^{-kt}$, and then k is the rate constant for the chemical reaction leading out of that well into the other state.

Problem 38: What’s the alternative? You should think a bit about what was just said. Suppose for example, that you don’t know the potential and I just give you samples of the trajectory $x(t)$. What would it mean if the trajectories paused at some intermediate point between reactants and products? How would you interpret non-exponential distributions of the time spent in each well?

Problem 39: Numerical experiments on activation over a barrier. Perhaps before launching into the long calculation that follows, you should get a feeling for the problem by doing a small simulation. Consider a particle at position x moving in a potential $V(x) = V_0[1 - (x/x_0)^2]^2$. Notice that this is double well, with minima at $x = \pm x_0$ and a barrier of height V_0 between these minima. Let’s consider the overdamped limit of Brownian motion in this potential, as in Eq (224),

$$\gamma \frac{dx(t)}{dt} = \frac{4V_0}{x_0} \left(\frac{x}{x_0} \right) \left[1 - \left(\frac{x}{x_0} \right)^2 \right] + \zeta(t), \quad (226)$$

We want to simulate these dynamics. The simplest approach is the naive one, in which we use discrete time steps separated by Δt and we approximate

$$\frac{dx(t)}{dt} \rightarrow \frac{x(n+1) - x(n)}{\Delta t}. \quad (227)$$

(a.) To use this discretization we have to deal with the Langevin force. One (moderately) systematic approach is to integrate the Langevin equation over a small window of time Δt :

$$\gamma \int_t^{t+\Delta t} dt \frac{dx(t)}{dt} = - \int_t^{t+\Delta t} dt \frac{\partial V(x)}{\partial x} + \int_t^{t+\Delta t} dt \delta F(t) \quad (228)$$

$$\gamma [x(t+\Delta t) - x(t)] \approx -\Delta t \left. \frac{\partial V(x)}{\partial x} \right|_{x=x(t)} + z(t), \quad (229)$$

where

$$z(t) = \int_t^{t+\Delta t} dt \zeta F(t). \quad (230)$$

Using the correlation function of the Langevin force from Eq (225), compute the variance of $z(t)$. Show also that the values of z at different times—separated at least by one discrete step Δt —are uncorrelated.

(b.) Combine your results in [a] with the equations above to show that this simple discretization is equivalent to

$$y(n+1) = y(n) + \alpha E^\ddagger \cdot y(n) \cdot [1 - y^2(n)] + \sqrt{\frac{\alpha}{2}} \zeta(n), \quad (231)$$

where $y = x/x_0$, the parameter $\alpha = 4k_B T \Delta t / (\gamma x_0^2)$ should be small, $E^\ddagger = V_0 / (k_B T)$ is the normalized “activation energy” for escape over the barrier, and $\zeta(n)$ is a Gaussian random number with zero mean, unit variance, and no correlations among different time steps n .

(c.) Implement Eq (231), for example in MATLAB. Note that MATLAB has a command `randn` that generates Gaussian random numbers.³² You might start with a small value of E^\ddagger , and experiment to see how small you need to make α before the results start to make sense. What do you check to see if α is small enough?




FIG. 33 Example of the trajectories we expect to see in solving the Langevin Eq (224). Long sojourns in the reactant or product state are interrupted by rapid jumps from one potential well to the other. If we look at the times t_r spent in the reactant state, these should come from a probability distribution $P_r(t_r) = k_+ e^{-k_+ t_r}$, where k_+ is the rate of the chemical reaction from reactants to products. Similarly we should have $P_p(t_p) = k_- e^{-k_- t_p}$, where k_- is the rate of the reverse reaction.

³² More precisely, MATLAB *claims* that `randn` generates Gaussian random numbers that are independent. Maybe you should check this?

(d.) Explore what happens as you change the value of E^\ddagger . For each value of E^\ddagger , check that your simulation runs long enough so that the distribution of x actually is given by the Boltzmann distribution, $P(x) \propto \exp[-V(x)/k_B T]$. As E^\ddagger increases, can you see that there are isolated discrete events corresponding to the “chemical reaction” in which the system jumps from one well to the other? Use your simulation to estimate the rate of these jumps, and plot the rate as a function of the activation energy E^\ddagger . Can you verify the Arrhenius law?

Problem 40: Effective potentials. We are discussing, for simplicity, a one dimensional problem. Suppose that there are really many dimensions, not just x but also $y_1, y_2, \dots, y_N \equiv \{y_j\}$. Then we have, again in the overdamped limit,

$$\gamma \frac{dx}{dt} = -\frac{\partial V(x; \{y_j\})}{\partial x} + \zeta(t) \quad (232)$$

$$\gamma_i \frac{dy_i}{dt} = -\frac{\partial V(x; \{y_j\})}{\partial y_i} + \xi_i(t), \quad (233)$$

where, as usual

$$\langle \xi_i(t) \xi_j(t') \rangle = 2k_B T \gamma_i \delta_{ij} \delta(t - t'). \quad (234)$$

Imagine now that x moves much more slowly than all the $\{y_j\}$.

(a.) Verify that, from Eq (233), the stationary distribution of $\{y_j\}$ at fixed x is the Boltzmann distribution,

$$P(\{y_j\}|x) = \frac{1}{Z(x)} \exp \left[-\frac{V(x; \{y_j\})}{k_B T} \right]. \quad (235)$$

(b.) If x is slow compared with all the $\{y_j\}$, it is plausible that we should average the dynamics of x in Eq (232) over the stationary distribution $P(\{y_j\}|x)$. Show that this generates an equation in which x moves in an effective potential,

$$\gamma \frac{dx}{dt} = -\frac{\partial V_{\text{eff}}(x)}{\partial x} + \zeta F(t), \quad (236)$$

and this effective potential is the free energy, $V_{\text{eff}}(x) = -k_B T \ln Z(x)$.

(c.) Equations (232) and (233) still aren't completely general, since we have taken the mobility tensor to be diagonal, so that forces on coordinate y_i lead to velocities only along this direction. Does the more general case presents any new difficulties for the problem posed here?

This picture of trajectories that hover around one well and then jump to another should remind you of something you learned in quantum mechanics. In particular, if you take the path integral view of quantum mechanics, then tunneling in a double well potential is dominated by these sorts of trajectories. In fact, if Planck's constant is small, so that tunneling is rare, there is a semi-classical approximation to the path integral which reproduces the WKB approximation to Schrödinger's equation, and in this approximation the path integral is dominated by specific trajectories, which have come to be called “instantons.” These instantons are precisely the jumps from one well to another, analogous to what we have drawn for the classical case in Fig 33.

There are three seemingly different but equivalent ways of doing quantum mechanics. Most elementary courses focus on Schrödinger's equation, which describes the amplitude for a particle to be at position x at time t . But you can also look at Heisenberg's equations of motion for

the position (and momentum) operator, and finally one can use path integrals. How do these different approaches to quantum mechanics connect with the description of Brownian motion?

The Langevin equation is a bit like Heisenberg's equation for the position operator. It seems to give us something most closely related to the equations of motion in classical (noiseless) mechanics, but it requires some interpretation. In the case of the Langevin equation, because $\zeta(t)$ is random, when we solve for the trajectory $x(t)$ we get something different for every realization of ζ , so “solve” should be used carefully. More precisely what we get, for example, from one simulation of the Langevin equation is a sample drawn out of the distribution of trajectories.

When we pass from Heisenberg's equations of motion to the Schrödinger equation, we shift from trying to follow the time dependence of coordinates to trying to see the whole distribution of coordinates at each time, as encoded in the wave function. Similarly, we can pass from the Langevin equation to the diffusion equation, which governs the probability $P(x, t)$ that we will find the particle at position x at time t . It is useful to remember that the diffusion equation is an equation for the conservation of probability,

$$\frac{\partial P(x, t)}{\partial t} = -\frac{\partial}{\partial x} J(x, t), \quad (237)$$

where $J(x)$ is the probability current.³³ Fick's law tells us that diffusion contributes a current that tends to reduce gradients in the concentration of particles, or equivalently gradients in the probability of finding one particle, so that

$$J_{\text{diff}}(x, t) = -D \frac{\partial P(x, t)}{\partial x}. \quad (238)$$

But if there is some force $F(x) = -dV(x)/dx$ acting on the particle, it will move with an average velocity $v = F(x)/\gamma$, and hence there is a ‘drift’ current

$$J_{\text{drift}}(x, t) = vP(x, t) = -\frac{1}{\gamma} \frac{dV(x)}{dx} P(x, t). \quad (239)$$

³³ A note about units. Often when discussing diffusion it is natural to think about the concentration of particles, which has units of particles per unit volume. The current of particles then has units of particles per area per time. What we are doing here is slightly different. First, we are talking about the probability of finding *one* particle at point x . Second, we are in one dimension, and so this probability distribution has units of 1/(length), not 1/(volume). Then the current has the units of a rate, 1/(time). Check that this make the units come out right in Eq's (??) and (238).

Putting these terms together, $J = J_{\text{diff}} + J_{\text{drift}}$, we have

$$\begin{aligned} \frac{\partial P(x,t)}{\partial t} &= -\frac{\partial}{\partial x} \left[-D \frac{\partial P(x,t)}{\partial x} - \frac{1}{\gamma} \frac{dV(x)}{dx} P(x,t) \right] \\ &= D \frac{\partial}{\partial x} \left[\frac{\partial P(x,t)}{\partial x} + \frac{1}{\gamma D} \frac{dV(x)}{dx} P(x,t) \right] \\ &= D \frac{\partial}{\partial x} \left[\frac{\partial P(x,t)}{\partial x} + \frac{1}{k_B T} \frac{dV(x)}{dx} P(x,t) \right], \end{aligned} \quad (240)$$

where in the last step we use the Einstein relation $D = k_B T / \gamma$. This way of writing the diffusion equation makes clear that the Boltzmann distribution $P \propto e^{-V(x)/k_B T}$ is an equilibrium ($\partial P / \partial t = 0$) solution.

We have said that, in looking at solutions of the Langevin equation, the signature of a ‘‘chemical reaction’’ with rate k is that the trajectories $x(t)$ will look like they do in Fig 33. What is the corresponding signature in the solutions of the diffusion equation? More precisely, even if we solve the diffusion equation to get the full $P(x,t)$ from some initial condition, what is it about this solution that corresponds to the rate constant k ? In the same way that Schrödinger’s equation is a linear equation for the wave function, the diffusion equation is a linear equation for the probability, which we can write as

$$\frac{\partial P(x,t)}{\partial t} = \hat{L}P(x,t). \quad (241)$$

All the dynamics are determined by the eigenvalues of the linear operator \hat{L} :

$$P(x,t) = \sum_n a_n e^{\lambda_n t} u_n(x) \quad (242)$$

$$\hat{L}u_n(x) = \lambda_n u_n(x). \quad (243)$$

We know that one of the eigenvalues has to be zero, since if $P(x,t)$ is the Boltzmann distribution, $P \propto e^{-V(x)/k_B T}$, it won’t change in time. Deviations from the Boltzmann distribution should decay in time, so all the nonzero eigenvalues should be negative.

Problem 41: Positive decay rates. We know that $P \propto \exp[-V(x)/k_B T]$ is a stationary solution of the diffusion Eq (240). To study the dynamics of how this equilibrium is approached, write

$$P(x,t) = \exp \left[-\frac{V(x)}{2k_B T} \right] Q(x,t). \quad (244)$$

(a.) Derive the equation governing $Q(x,t)$. Show that (by introducing factors of i in the right place) this can be written as

$$\frac{\partial Q(x,t)}{\partial t} = -A^\dagger A Q(x,t), \quad (245)$$

where the combination $A^\dagger A$ is a Hermitian operator. This gives an explicit version of Eq (241); explain why this implies that all the eigenvalues $\lambda_n \leq 0$.

(b.) For the case of the harmonic potential, $V(x) = \kappa x^2/2$, show that the operators A^\dagger and A are the familiar creation and

annihilation operators from the quantum harmonic oscillator. Use this mapping to find all the eigenvalues λ_n . How do these relate to the time constant for exponential decay that you get from the noiseless dynamics [Eq (224) with $\zeta(t) = 0$]?

If we place the molecule in some configuration that is far from the local minima in each potential well, it will ‘slide’ relatively quickly into its relaxed configuration, and execute some Brownian motion around this sliding trajectory so that it samples the Boltzmann distribution within the well. This relaxation should be described by some of the eigenvalues λ_n , and these should be large and negative, corresponding to fast relaxation. In practice, we know that molecules in solution achieve this sort of ‘vibrational relaxation’ within nanoseconds if not picoseconds.

The statement that there is a chemical reaction at rate k means that, as a population of molecules comes to equilibrium, all the equilibration *within* the reactant or product states is fast, corresponding to time scales much shorter than $1/k$. On the much longer time scale $1/k$, there is equilibration between the reactant and product states. Thus, if we look at the whole spectrum of eigenvalues λ_n for the diffusion equation, one eigenvalue should be zero (as noted above), almost all the others should be very large and negative, while there should be one isolated eigenvalue that is small and negative—and this will be the reaction rate k , or more precisely the sum of the rates for the forward (reactants \rightarrow products) and backward (products \rightarrow reactants) reactions.




FIG. 34 Decay rates in diffusion compared with energy levels in quantum mechanics. In both cases there is a small splitting between the first two eigenvalues. For the diffusive case, this splitting is the rate of thermally activated hopping over the barrier—a chemical reaction. For the quantum case this splitting is the tunneling frequency between the two wells.

We arrive, then, at a picture of the eigenvalue spectrum in which there is a small splitting (between $\lambda_0 = 0$ and $\lambda_1 = -k$) relative to the next highest eigenvalue, as shown in Fig 34. This should remind you of what happens in quantum mechanical tunneling between two potential wells. The basic spacing of energy levels is set by the vibrational quanta within each well, but the states—and, in particular, the ground state—is split by a small amount corresponding to the frequency of tunneling between the two wells. It is the size of the barrier, or equivalently the smallness of \hbar , which makes this splitting small. In the diffusion problem, it is presumably the smallness of the temperature relative to the activation energy which enforces $\lambda_0 - \lambda_1 \ll \lambda_1 - \lambda_2$. We know how to solve Schrödinger's equation using the WKB approximation to extract the small tunneling amplitude, and so there should be a similar approximation to the diffusion equation that allows us to calculate the reaction rate.

The WKB approximation has a natural formulation in the path integral approach—in the limit $\hbar \rightarrow 0$, the path integral describing the amplitude for any quantum process is dominated by particular trajectories that are solutions of the classical equations of motion, although for classically forbidden processes (as with tunneling) these equations have to be continued to imaginary time. This idea of a dominant trajectory should be even clearer in the case of Brownian motion, since we won't have to deal with the continuation to imaginary time. To see how this works—and, finally, to derive the Arrhenius law—we need to construct the probability distribution functional for the trajectories $x(t)$ that solve the Langevin Eq (224).

The probability that we observe a trajectory $x(t)$ can be calculated by finding the random force $\zeta(t)$ which was needed to generate this trajectory, and then calculating the probability of this force. We know that the random forces come from a Gaussian distribution, and we know the correlation function [Eq (225)], so we have

$$P[\zeta(t)] \propto \exp \left[-\frac{1}{4\gamma k_B T} \int dt \zeta^2(t) \right]. \quad (246)$$

The Langevin equation, Eq (224), can be rewritten as

$$\zeta(t) = \gamma \frac{dx}{dt} + \frac{dV(x)}{dx}, \quad (247)$$

so it is tempting to say that the probability of observing the trajectory $x(t)$ is given by

$$P[x(t)] \sim \exp \left[-\frac{1}{4\gamma k_B T} \int dt \left(\gamma \frac{dx}{dt} + \frac{dV(x)}{dx} \right)^2 \right], \quad (248)$$

and this is almost correct. To see what's missing, consider the simpler case where we just have one variable x [instead of a function $x(t)$] that obeys an equation

$$f(x) = y, \quad (249)$$

and y is random, drawn from a distribution $P_y(y)$. It is tempting to write

$$P_x(x) = P_y(y = f(x)), \quad (250)$$

but this can't be right— x and y can have different units, and hence P_x and P_y must have different units. As you have probably seen many times before, in this simple one dimensional example, the correct statement is that the probability mass within some small region dx must be equal to the mass found in the corresponding dy ,

$$P_x(x)dx = P_y(y = f(x))dy \quad (251)$$

$$\Rightarrow P_x(x) = P_y(y = f(x)) \left| \frac{dy}{dx} \right| \quad (252)$$

$$= P_y(y = f(x)) \left| \frac{df(x)}{dx} \right|. \quad (253)$$

More generally, in order to equate probability distributions, we need a Jacobian for the transformation between variables. Thus, instead of Eq (248), we really want to write

$$P[x(t)] \propto \exp \left[-\frac{1}{4\gamma k_B T} \int dt \left(\gamma \frac{dx}{dt} + \frac{dV(x)}{dx} \right)^2 \right] \mathcal{J}, \quad (254)$$

where \mathcal{J} is the Jacobian of the transformation between $x(t)$ and $\delta F(t)$. Importantly, the Jacobian doesn't depend on temperature. In contrast, the exponential term that we have written out is $\sim e^{-1/T}$, so at low temperatures this will dominate. So, for this discussion, we won't worry about the Jacobian.

Problem 42: Jacobians. [Give a problem that walks through the derivation of the Jacobian, as in Zinn-Justin.]

To make use of Eq (254), it's useful to look more closely at the integral which appears in the exponential. Let's be careful to let time run from some initial time t_i up to some final time t_f :

$$\int_{t_i}^{t_f} dt \left(\gamma \frac{dx}{dt} + \frac{dV(x)}{dx} \right)^2 = \int_{t_i}^{t_f} dt \left[\left(\gamma \frac{dx}{dt} \right)^2 + 2\gamma \frac{dx}{dt} \frac{dV(x)}{dx} + \left(\frac{dV(x)}{dx} \right)^2 \right] \quad (255)$$

$$= \int_{t_i}^{t_f} dt \left[\left(\gamma \frac{dx}{dt} \right)^2 + \left(\frac{dV(x)}{dx} \right)^2 \right] + 2\gamma \int_{t_i}^{t_f} dt \frac{dx}{dt} \frac{dV(x)}{dx} \quad (256)$$

$$= \int_{t_i}^{t_f} dt \left[\left(\gamma \frac{dx}{dt} \right)^2 + \left(\frac{dV(x)}{dx} \right)^2 \right] + 2\gamma \int_{t_i}^{t_f} dt \frac{dV(x)}{dt}, \quad (257)$$

$$= \int_{t_i}^{t_f} dt \left[\left(\gamma \frac{dx}{dt} \right)^2 + \left(\frac{dV(x)}{dx} \right)^2 \right] + 2\gamma[V(x_f) - V(x_i)], \quad (258)$$

where in the last steps we recognize one term as a total derivative; as usual $x_i = x(t_i)$ is the initial position, and similarly $x_f = x(t_f)$ is the final position. Substituting, we can write the probability of a trajectory $x(t)$ as

$$P[x(t)] \propto \mathcal{J} e^{-S/k_B T}, \quad (259)$$

where the ‘action’ takes the form

$$S = \frac{V(x_f) - V(x_i)}{2} + \int_{t_i}^{t_f} dt \left[\frac{\gamma}{4} \left(\frac{dx}{dt} \right)^2 + \frac{1}{4\gamma} \left(\frac{dV(x)}{dx} \right)^2 \right]. \quad (260)$$

This is a good time to remember that, for the simplest problems of classical mechanics, the action is

$$S_{\text{cm}} = \int_{t_i}^{t_f} dt \left[\frac{m}{2} \left(\frac{dx}{dt} \right)^2 - \mathcal{U}(x(t)) \right], \quad (261)$$

where m is the mass and $\mathcal{U}(x)$ is the potential energy. Except for a constant, the effective action for our problem is exactly that of a simple mechanics problem of a particle with mass m moving in a effective potential $\mathcal{U}(x)$,

$$m = \frac{\gamma}{2} \quad (262)$$

$$\mathcal{U}(x) = -\frac{1}{4\gamma} \left(\frac{dV(x)}{dx} \right)^2. \quad (263)$$

Figure 35 shows how this effective potential relates to the original double well.

At low temperatures, the distribution of trajectories will be dominated by those which minimize the action S . Clearly, one way to make the action minimal (zero, in fact) is to have the position be constant at one of the minima of the potential well. This describes a situation in which nothing happens. To have a chemical reaction, we need a trajectory that starts in the well corresponding to the reactants state, climbs up to the ‘transition state’ at the top of the barrier, and then slides down the other side. Let’s start with the first part of this problem, finding a trajectory that climbs the barrier. The dominant trajectory of this form will be one that minimizes the action, and from Fig 35 we see that this is equivalent to finding the solution to an ordinary mechanics problem

in which a particle starts on top of one hill, slides down and then gently comes to rest at the top of the next hill.

Problem 43: Zero energy? What we have just described are trajectories in the effective potential that have zero energy. There are, of course, trajectories that minimize the action but have nonzero energy. Why don’t we consider these?

Taking the details of Fig 35 seriously, if we start at rest on top of one hill, this means that we start with zero energy. But energy is conserved along the trajectory, so that

$$\frac{m}{2} \left(\frac{dx}{dt} \right)^2 + \mathcal{U}(x(t)) = E = 0. \quad (264)$$




FIG. 35 Potentials and forces in a double well. Top panel show the true potential, middle panel the force, and bottom panel the effective potential that enters the probability distribution of trajectories. Notice that each extremum of the potential, both maxima and minima, becomes a maximum of the effective potential, and all these maxima are degenerate at $\mathcal{U} = 0$.

This means that

$$\mathcal{U}(x(t)) = -\frac{m}{2} \left(\frac{dx}{dt} \right)^2 \quad (265)$$

$$\frac{dx}{dt} = \pm \sqrt{-\frac{2}{m} \mathcal{U}(x(t))}; \quad (266)$$

we are interested in trajectories that move from left to right, so we should choose the upper sign, so that $dx/dt > 0$. But now we can substitute into the action,

$$\begin{aligned} S_{\text{cm}} &= \int_{t_i}^{t_f} dt \left[\frac{m}{2} \left(\frac{dx}{dt} \right)^2 - \mathcal{U}(x(t)) \right] \\ &= \int_{t_i}^{t_f} dt \left[\frac{m}{2} \left(\frac{dx}{dt} \right)^2 + \frac{m}{2} \left(\frac{dx}{dt} \right)^2 \right] \end{aligned} \quad (267)$$

$$= m \int_{t_i}^{t_f} dt \left(\frac{dx}{dt} \right)^2 \quad (268)$$

$$= m \int_{t_i}^{t_f} dt \frac{dx}{dt} \sqrt{-\frac{2}{m} \mathcal{U}(x(t))}, \quad (269)$$

so that finally we have

$$S_{\text{cm}} = \int_{x_i}^{x_f} dx \sqrt{-2m\mathcal{U}(x)}, \quad (270)$$

where you should recognize the connection to the WKB formula for tunneling. In our case, the effective potential and mass are defined by Eq's (262) and (263), so that

$$-2m\mathcal{U}(x) = -2\frac{\gamma}{2} \left[-\frac{1}{4\gamma} \left(\frac{dV(x)}{dx} \right)^2 \right] = \frac{1}{4} \left(\frac{dV(x)}{dx} \right)^2. \quad (271)$$

It is quite nice how the factors of γ cancel. Substituting into Eq (270), we find

$$S_{\text{cm}} = \frac{1}{2} \int_{x_i}^{x_f} dx \sqrt{\left(\frac{dV(x)}{dx} \right)^2} \quad (272)$$

$$= \pm \frac{1}{2} \int_{x_i}^{x_f} dx \frac{dV(x)}{dx} \quad (273)$$

$$= \frac{1}{2} \left| V(x_f) - V(x_i) \right|, \quad (274)$$

where we choose the sign in taking the root so that the action comes out positive, as it must from Eq (268).

Problem 44: Extracting the dominant paths. We have seen that, in the low temperature limit, the reaction is dominated by trajectories that lead from one well to the other and minimize the action. Look through your simulation results from Problem 28, and collect as many examples as you can of the 'jumping' trajectories. How do these examples compare with the theoretical prediction that comes from minimizing the action? Can you align the sample trajectories well enough to compute an average that might be more directly comparable to the theory?

To finish the calculation, we need to put some of these pieces together. The action that determines the probability of a trajectory is, from Eq (260),

$$\begin{aligned} S &= \frac{V(x_f) - V(x_i)}{2} + \int_{t_i}^{t_f} dt \left[\frac{\gamma}{4} \left(\frac{dx}{dt} \right)^2 + \frac{1}{4\gamma} \left(\frac{dV(x)}{dx} \right)^2 \right] \\ &= \frac{V(x_f) - V(x_i)}{2} + S_{\text{cm}} \end{aligned} \quad (275)$$

$$= \frac{V(x_f) - V(x_i)}{2} + \frac{1}{2} \left| V(x_f) - V(x_i) \right|. \quad (276)$$

This is a remarkably simple result. If we are looking at a trajectory that climbs from the bottom of a potential well to the top of the barrier, we have $V(x_f) > V(x_i)$ and hence the action is

$$S_{\text{climb}} = V(x_f) - V(x_i) = E_{\text{act}}, \quad (277)$$

which is the "activation energy" for going over the barrier. On the other hand, if we look at a trajectory that slides down from the barrier into the other well, we have

$V(x_f) < V(x_i)$ and hence

$$S_{\text{slide}} = 0. \quad (278)$$

So, what we have shown is that paths which take us from reactants to products, climbing the barrier and sliding down the other side, have a minimal action $S_{\text{react}} = S_{\text{climb}} + S_{\text{slide}} = E_{\text{act}}$. Thus, the probability of seeing such a trajectory is

$$P[x_{\text{react}}(t)] \propto \mathcal{J} e^{-S_{\text{react}}/k_B T} \sim e^{-E_{\text{act}}/k_B T}, \quad (279)$$

and this is the essence of the Arrhenius law (at last).

One could legitimately complain that we haven't really solved our problem. All we have done is to show that, in some window of time, trajectories that jump from reactants to products are suppressed in probability by a factor $e^{-E_{\text{act}}/k_B T}$. This is the basic idea of the Arrhenius law, but we haven't actually calculated a rate constant. In truth, this last step requires rather more technical apparatus, in the same way that getting the tunneling rate in the WKB approximation is harder than getting the exponential suppression, so I will leave it aside for now.

So far, we have given a fairly general discussion, and perhaps it's not obvious whether there is anything special about how these ideas will play out in the case of biological molecules. If we try to draw the picture in Fig 32, we usually associate the "reaction coordinate," that is the molecular coordinate along which we see the double well potential, with the motions that are involved in the chemical events themselves. Thus, if we are looking at the transfer of a hydrogen atom, breaking one bond and forming another, we might think that the relevant molecular coordinate is given by the position of the hydrogen atom itself.

Biological molecules—such as the proteins which act as enzymes, catalyzing specific chemical reactions of importance in the cell—are large, and hence flexible. Certainly they can change reaction rates by holding the reactants in place. But because of their flexibility, there is also the possibility that, as they flex, the effective barrier for the reaction changes. In this case, the dominant path for the reaction might be for the protein to fluctuate into a favorable configuration, and then for the more local coordinates (e.g., the position of the hydrogen atom) to make their jump. In this way, the observed activation energy comes to have two components, the usual one that we measure along the reaction coordinate, which presumably is reduced by waiting for the protein to arrange itself properly, and then the energy of distorting the protein itself.

To be a little more formal, imagine that for every configuration Q of the protein, there is a different activation energy for the reaction, $E_{\text{act}}(Q)$. Of course there must also be some (free) energy of the protein once it is in the structure described by Q , and this determines the probability distribution $P(Q)$. Then if the fluctuations in Q are fast, we expect to see an average rate constant

$$k = A \int dQ P(Q) \exp \left[-\frac{E_{\text{act}}(Q)}{k_B T} \right]. \quad (280)$$

If we fix Q at its equilibrium position, we could find that $E_{\text{act}}(Q = Q_{\text{eq}})$ is large, which might make us think that the reaction will be slow. But by sampling non-equilibrium configurations, the protein can speed up the reaction.

Obviously this general picture depends on many details, but before proceeding one could ask if there is any

evidence for such coupling of protein structural fluctuations to the modulations of chemical reaction rates. I think the strongest evidence is from the mid 1970s, in a beautiful series of experiments by Austin and colleagues. The idea is very simple. Suppose that we really do have the activation energy varying with the configuration of the protein. If we could stop the motion of the protein, then each molecule would be stuck with a different activation energy and hence a different reaction rate. Then, instead of seeing an average rate, each molecule reacts at its own rate, and if we count the total number of molecules that have not yet reacted we should see

$$N(t) = \int dQ P(Q) \exp \left[-A e^{-E_{\text{act}}(Q)/k_B T} t \right], \quad (281)$$

which definitely is not an exponential decay. In fact if the fluctuations in Q generate very large variations in the activation energy, then this is very far from being an exponential decay.

Problem 45: Power law decays. Suppose that the effect of the fluctuations in Q is to generate a distribution of activation energies

$$P(E_{\text{act}}) = \frac{1}{n! E_0^n} (E/E_0)^n e^{-E/E_0}. \quad (282)$$

Then we should have

$$N(t) = \int_0^\infty \frac{dE}{n! E_0^n} (E/E_0)^n e^{-E/E_0} \exp \left[-A e^{-E/k_B T} t \right]. \quad (283)$$

(a.) Show that, at large t , there is a saddle point approximation to this integral, and that this predicts a decay $N(t) \sim t^{-\alpha}$. What determines the power α ? Are there corrections to this formula?

(b.) Calculate the average rate constant, as in Eq (280),

$$\bar{k} = A \int_0^\infty \frac{dE}{n! E_0^n} (E/E_0)^n e^{-E/E_0} \exp \left[-\frac{E}{k_B T} \right]. \quad (284)$$

Does this mean rate obey the Arrhenius law? How large are the deviations? Is there a limit in which the Arrhenius law is recovered?

Problem 46: A more sophisticated model. The discussion above concerns either the limit in which fluctuations are very fast, so we see an average rate constant, or very slow, so that we see a distribution of rate constants. It would be nice to have a simple model that interpolates between these limits. [\[give a problem that goes through the essence of the Agmon & Hopfield papers ...\]](#)

So, we have the dramatic prediction that if we would freeze the motion of the protein, we'd see something very far from the usual exponential decays. In order to test this we need the right model system. In particular, if we are literally going to freeze things, then molecules can't diffuse relative to one another, and most what we usually think of as chemistry will stop. We need an example of a reaction that happens among molecules that are already "together" and ready to react. If things are frozen, then




FIG. 36 The heme group at the center of myoglobin, hemoglobin, and other heme proteins. Recall the convention (Fig 15) that carbon atoms are at unmarked nodes of the skeleton, and hydrogen atoms which complete the four bonds needed for each carbon are not shown. The iron atom at the center is also coordinated from below by a nitrogen from the protein, and oxygen or carbon monoxide can bind to the iron from above the plane. The large conjugated structure of the heme group endows the molecule with a strong absorption cross section in the visible and ultraviolet range of the spectrum. Because the electronic states of the heme mix with the d orbitals of the iron, the absorption spectrum shifts upon oxygen binding.

the usual trick of suddenly mixing the reactants together to start the reaction also isn't going to work.

In many organisms, including us, oxygen is essential for a wide variety of processes. We take in oxygen by breathing, and need to distribute it to all of our tissues. The way we do this is to have specific proteins to which oxygen binds, and then the proteins are transported, starting in the blood. The major such oxygen transport protein in our blood is called hemoglobin, which is described in more detail in Appendix A.4 because it provides the classic example of cooperativity in protein function. Hemoglobin has four protein subunits, each of which can bind a single oxygen molecule. In our muscles we find a simpler protein, with just one subunit, called myoglobin. Myoglobin, hemoglobin, and the cytochromes that we will discuss below all are members of the “heme protein” family, which are defined by the fact that they bind a rather large planar organic molecule called heme, with an iron atom at its center, as shown in Fig 36. This iron is held in the plane by nitrogens from the heme and from below by a nitrogen from the protein. Oxygen can bind to the iron from above the plane.

The iron atom, and hence the oxygen binding site is buried deep inside the protein, as shown in Fig 37. This is interesting in part because it tells us that the full process of binding and unbinding must involve some motion

or “breathing” of the protein structure. Further, once oxygen binds, if we freeze the protein it will be trapped, unable to escape. The conjugated electronic structure of the heme generates a strong optical absorption band, and because the electronic states of the heme mix with the orbitals of the iron, the absorption shifts when oxygen binds to the iron. Further, when a photon is absorbed by myoglobin with oxygen bound, there is some probability that the energy of the absorbed photon will be channeled into breaking the bond between the iron and the oxygen. Thus, if we let oxygen bind to myoglobin and then freeze the solution, we can knock the oxygen off the iron atom with a flash of light, and then we can watch the oxygen rebound after rattling around in the “pocket” formed by the protein.

In principle, motion of the oxygen molecule from the pocket to the iron atom needn't be coupled to motions of the protein. But if this coupling does occur, we expect, from the discussion above, that the kinetics of the rebinding after a light flash will deviate strongly from an exponential decay. We can follow the kinetics by looking at the absorption spectrum, and this is what is shown in Fig ?? for both oxygen and carbon monoxide binding to myoglobin. We see that once the solution is truly frozen solid (below ~ 160 K in the glycerol–water mixtures used for these experiments), the fraction of molecules that have not reacted decays more nearly as a power law than an exponential. This suggests that we have frozen in a very




FIG. 37 A slice through the electron density map of the myoglobin molecule, as inferred from X-ray diffraction data (Kendrew 1964). This map is made from data at 1.4 \AA resolution. In the center we see the heme group edge on. The histidine side chain from the protein coordinates the iron atom from below the plane of the heme, and in the crystals used in these experiments a water molecule binds to the iron atom in the position that would be taken by oxygen. Note that there is not much empty space in the structure, so that the protein actually has to “breathe” in order for oxygen to have access to the iron, or to escape once bound.

broad distribution of rate constants.




FIG. 38 Rebinding of oxygen and carbon monoxide to myoglobin at low temperatures, following a flash of light to break the bond (Austin et al 1975). Circle are data points, obtained by monitoring the absorption spectrum. Note that this is a logarithmic plot on *both* axes, so that we see an enormous range of times. Lines are fits to the phenomenological power law decay $N(t) = 1/[1 + (t/t_0)]^n$. The dashed line shows, for contrast, an exponential decay, $N(t) = e^{-kt}$, with $k = 1 \text{ s}^{-1}$.

So far our discussion of chemical reactions has treated motion along the reaction coordinate as being completely classical. Is it possible that quantum effects could be relevant? Notice in Fig 38 that as we lower the temperature, the kinetics remain consistently non-exponential, but the typical time scale (e.g., the time required for the reach to reach 90% completion) is slowing down. If we keep lowering the temperature, eventually this slowing stops, and we see temperature independent kinetics. Almost certainly this arises because the reaction proceed by quantum mechanical tunneling through the effective barrier rather than by thermal activation over the barrier. The observation of quantum mechanical effects in a biological system always triggers excitement, although this is tempered somewhat by the fact that, in this case, to see tunneling one has to go to very low temperatures (below 10 K) indeed. In fact, well before the work on myoglobin, there had been observations of temperature independent kinetics in the photon-triggered electron transfer reactions in photosynthesis. Although our immediate experience of photosynthesis involved plants, many of the key experiments on the dynamics of electron transfer were done in photosynthetic bacteria.

The basic business of photosynthesis is carried out by the reaction center, a huge complex of proteins that holds a collection of medium sized, organic molecules—chlorophylls, pheophytins (chlorophylls without the magnesium), and quinones. [Need some schematics of these molecules, plus the reaction center structure.] Two of the chlorophylls are held in a special pair (P), and the elec-

tronic states of these two molecules are strongly mixed. If one purifies the reaction center away from all the accessory structures, the photochemistry is triggered when the special pair absorbs a photon.

From the excited state of the special pair, an electron hops to states localized on the pheophytin (I) and then the quinone (Q), as shown in Fig 39. Because P and Q are held, by the protein scaffold, on opposite sides of a membrane, the net effect is to transfer charge across the membrane, capturing the energy of the absorbed photon. Quinones [point back to the structure!] exist in multiple protonation states, so that the electron transfer can couple to proton transfer, and in this way the reaction center serves to drive protons across the membrane. The difference in electrochemical potential for protons provides a universal energy source that is used by other transmembrane protein, for example to synthesize ATP, which all cells use in powering other processes (including movement). In more complex organisms, including green plants, there are two kinds of reaction centers, one of which couples photon-driven electron transfer to the splitting of water to make all the oxygen in our atmosphere.

To complete the cycle and “reset” the reaction center for the arrival of the next photon, the hole on the special pair needs to be filled in, and this happens by electron transfer from another protein, cytochrome c, which can also diffuse away from the membrane and interact with the rest of the cell’s chemistry. It is this reaction that




FIG. 39 Schematic of the electron transfer reactions in the reaction center of photosynthetic bacteria. The “pigment molecule” P absorbs a photon, and transfers an electron from the excited state to an intermediate acceptor I, which then passes the electron to a quinone molecule Q; there is a second quinone, not shown. The hole on P is filled in by electron transfer from another protein, cytochrome c, and the kinetics of the reaction $CP^+ \rightarrow C^+P$ provided the first evidence for quantum tunneling in a biological system, as shown (redrawn from DeVault & Chance 1966).

provided the first evidence for tunneling in a biological system. If the cytochrome is absent, as in purified reaction centers, one can observe the recombination reaction $P^+Q^- \rightarrow PQ$, which also has an anomalous temperature dependence, as discussed below. To connect with the discussion of myoglobin, this recombination reaction also exhibits non-exponential kinetics under some conditions, suggesting that it is possible to freeze some of the fluctuations in structure that normally provide rapid modulations of the reaction rate.

The key to experiments on the kinetics of photosynthetic electron transfer is that all of the molecules involved change their absorption spectra significantly when they gain or lose an electron; not coincidentally, these spectra overlap with the spectrum of the solar radiation, and are concentrated in a range of wavelengths surrounding the ‘visible,’ from the near infrared to the near ultraviolet. We can trigger the reactions with a pulsed laser tuned to the absorption band of P, and we can then monitor different spectral features that track the different components. This started in the 1950s and 60s with time resolution in the microsecond range, and evolved—with successive revolutions in the techniques for generating short laser pulses—down to picoseconds and femtoseconds; this development parallels the exploration of the visual pigments described in Section I.B.

One key point about the photosynthetic reaction center is that all the electron transfer processes work even when the system is frozen, which tells us that there is no need for the different components to diffuse in order to find one another—all of the donor and acceptor sites are held in place by the protein scaffolding. This allows for investigation of the electron transfer reactions over a wide range of temperatures, and this was done to dramatic effect by DeVault and Chance in the mid 1960s, with the result shown in Fig 39. Near room temperature, the electron transfer from cytochrome c back to the special pair exhibits a normal, Arrhenius temperature dependence with an activation energy $E_{act} \sim 0.18$ eV. Importantly, the temperature dependence is continuous as the system is cooled through the solvent’s freezing point. But somewhere around $T \sim 100$ K, the temperature dependence stops, and the reaction rate remains the same down to liquid helium temperatures ($T \sim 4$ K). This strongly suggests that the reaction proceeds by tunneling at low temperatures.

Problem 47: A wrong model. If a reaction proceeds by activation over a barrier of height E , the rate is $k \propto \exp(-E/k_B T)$. If it proceeds by tunneling through the barrier, we expect $k \propto \exp(-2\sqrt{2mE}\ell/\hbar)$, where ℓ is the width of the barrier and m is the effective mass of the tunneling particle. For the DeVault–Chance reaction, there is a direct measurement of the activation energy $E \sim 0.18$ eV. If you imagine that it is the electron which has to go over or through this barrier, what value of ℓ is needed to

explain that the crossover from Arrhenius behavior to temperature independence occurs near $T \sim 100$ K? Does this result make any sense?

After roughly a decade of confusion (including discussions of the model in the previous problem), a clearer understanding of tunneling in electron transfer emerged in the mid to late 1970s.³⁴ The basic idea is schematized in Fig 40. We have an electron donor D and an acceptor A; the reaction is $DA \rightarrow D^+A^-$. The states DA and D^+A^- are different *electronic* states of the system. From the Born–Oppenheimer approximation, we know that when a molecule shifts to a new electronic state, the nuclei move on a new potential surface. We usually describe these nuclear or atomic motions as molecular vibrations, so we’ll refer to the relevant coordinates as vibrational coordinates. The simplest scheme, as in Fig 40, is one in which the vibrations are approximately harmonic. Then when we change electronic states, we can imagine changes in the structure of the normal modes, changes in the frequencies of these modes, and shifts in the equilibrium positions along the modes; barring symmetries, the last effect should be the leading one, and certainly it is the simplest.

In the state DA, an electron is localized on the donor. In the state D^+A^- , this electron is localized on the acceptor. If the donor and acceptor sites are far apart, as is often the case in large, biological molecules, then the




FIG. 40 Electron transfer is coupled to vibrational motion. ...

³⁴ The relevant physics here is essentially the same as in the discussion of absorption spectra in large molecules. See Chapter One and the Appendix on electronic transitions in large molecules; [give pointer to arXiv version](#).

wave functions of the electrons in these localized state will overlap only deep in their tails; any matrix element that connects these two states then must be very small. But if we want to have a transition between two states that are connected by only a small matrix element, then by Fermi's golden rule we need these states to be of the same energy. As shown in Fig 40, this happens only at special points, where the two potential energy surfaces for vibrational motion cross. The rate of the reaction should then be proportional to the probability of finding the system at this crossing point. The key point, then, is that at high temperatures this probability is controlled by the thermal fluctuations in the vibrational coordinates, while at low temperatures the system can still reach the crossing point, but now the fluctuations are dominated by quantum zero-point motion. If the activation energy—the energy required to distort the molecule from its equilibrium structure in state DA to the crossing point—is large compared to the relevant vibrational quanta, then

a zero-point fluctuation that carries the system to the crossing point necessarily involves sampling the tails of the ground state wavefunction, and this means that the system moves into a region that would be forbidden to a classical particle, even granting that it has the zero-point energy to work with. Thus, at low temperatures, the reaction is controlled by tunneling of the vibrational degrees of freedom, while at high temperatures these degrees of freedom move classically over the barrier.

To make all this a bit more precise, let's write the Hamiltonian corresponding to Fig 40. We have two electronic states, which we can take as the up and down states of a spin one-half. There is an energy difference between these states, which we'll call ϵ , and a weak matrix element Δ that mixes these states. There is a vibrational coordinate Q , and this coordinate moves in a potential that depends on the electronic state. Thus we have

$$\mathbf{H} = \frac{\epsilon}{2}\sigma_z + \Delta\sigma_x + \frac{1}{2}\dot{Q}^2 + \frac{1+\sigma_z}{2}V_\uparrow(Q) + \frac{1-\sigma_z}{2}V_\downarrow(Q), \quad (285)$$

If we think semi-classically, then the vibrational coordinates move hardly at all during the electronic transition, and so from the golden rule we should have the reaction rate

$$k \sim \frac{1}{\hbar}\Delta^2 \left\langle \delta(E_\uparrow - E_\downarrow) \right\rangle = \frac{1}{\hbar}\Delta^2 \left\langle \delta[\epsilon + V_\uparrow(Q) - V_\downarrow(Q)] \right\rangle, \quad (286)$$

where we have to average over the fluctuations of Q in the initial state DA. In the simplest case, where the potential surfaces are harmonic, differing only in their equilibrium positions,

$$V_\uparrow(Q) = \frac{\kappa}{2}Q^2 \quad (287)$$

$$V_\downarrow(Q) = \frac{\kappa}{2}(Q - Q_0)^2, \quad (288)$$

and hence $V_\uparrow(Q) - V_\downarrow(Q) = \kappa(Q_0Q - Q_0^2/2)$, so that

$$k \sim \frac{1}{\hbar}\Delta^2 \left\langle \delta\left(\epsilon - \frac{\kappa}{2}Q_0^2 + \kappa Q_0Q\right) \right\rangle \quad (289)$$

$$= \frac{\Delta^2}{\hbar\kappa Q_0} P\left(Q = \frac{Q_0}{2} - \frac{\epsilon}{\kappa Q_0}\right). \quad (290)$$

If we have a particle moving in a harmonic potential with frequency ω , then in thermal equilibrium the distribution of Q is Gaussian. The variance is $\langle(\delta Q)^2\rangle = k_B T_{\text{eff}}/\kappa$, where

$$k_B T_{\text{eff}} = \hbar\omega \left[\frac{1}{2} + \frac{1}{e^{\hbar\omega/k_B T} - 1} \right]; \quad (291)$$

notice that as $T \rightarrow 0$, $k_B T_{\text{eff}}$ approaches the zero-point

energy $\hbar\omega/2$. Putting all the terms together, we find

$$k \sim \frac{\Delta^2}{\hbar\sqrt{4\pi\lambda k_B T_{\text{eff}}}} \exp\left[-\frac{(\epsilon - \lambda)^2}{4\lambda k_B T_{\text{eff}}}\right], \quad (292)$$

where $\lambda = \kappa Q_0^2/2$ is the “reorganization energy” that would be required to distort the molecule from its equilibrium configuration in DA into the equilibrium configuration appropriate to D^+A^- if we didn't actually transfer the electron.

In Figure 41 we see the predicted dependence of the electron transfer rate on temperature in a parameter regime chosen to match the DeVault–Chance reaction. In order to have the transition between Arrhenius and tunneling behavior at the right temperature, we need a vibrational frequency $\omega/2\pi \sim 200 \text{ cm}^{-1}$.³⁵ If we look at the Raman spectra of cytochrome c or related molecules,

³⁵ Molecular vibrations contribute to the absorption of radiation in the infrared, and it is conventional to measure frequency in “wavenumbers” or inverse cm. To convert to the more usual Hz, just multiply by the speed of light, $3 \times 10^{10} \text{ cm/s}$. Note that this is a convention about units, and *not* a reference to the inverse wavelength in the medium used for the experiment, so




FIG. 41 Temperature dependence of the electron transfer rate, from Eq (292). Parameters are chosen, as described in the text, to match the behavior of the DeVault–Chance reaction in Fig 39. Circles are values of the rate computed at 20 K intervals, and dashed lines indicate the asymptotic behavior at high (activated) and low (tunneling) temperatures.

there is a vibrational mode near this frequency that corresponds to motions of the iron atom perpendicular to the plane of the heme group [obviously need a structural schematic!]. This makes sense, since when we add or subtract an electron from the molecule, this charge is shared between the iron and the heme, and on average the iron is displaced relative to the heme when the molecule changes its oxidation state. The energy difference between reactants and products can be measured directly by separate electrochemical experiments, and then to get the activation energy right we must have $\lambda \sim 0.14$ eV. If the relevant vibrational mode really is (mostly) the motion of the iron relative to the rest of the protein, then we know the mass associated the mode and hence the stiffness $\kappa = m\omega^2$, so we can determine $Q_0 \sim 0.2$ Å, and this is consistent with the displacements found upon comparing the oxidized and reduced structures of cytochrome c. So, this account of vibrational motion as controlling the temperature dependence of the reaction rate seems to make sense in light of everything else we know about these molecules, although admittedly it is a rough comparison. [Say something about the charge transfer band

there is no correction for the index of refraction. Once you start reading about molecular spectroscopy and chemical reactions (replete with calories and moles), you'll have to get some practice at changing units!

as a direct test?]

Problem 48: Getting numbers out. Convince yourself that the numbers in the preceding paragraph make sense. In particular, extract the estimate $Q_0 \sim 0.2$ Å for the motion of the iron atom relative to the protein.

There are many loose ends here. To begin, we have given a description in terms of one vibrational mode, but we have found an expression for the reaction rate that shows no sign of resonances when the energy difference ϵ between reactants and products in an integer multiple of the vibrational quantum $\hbar\omega$. Presumably the solution to the problem is the same as in our discussion of the absorption spectra of rhodopsin: individual modes are damped, so that resonances are broadened, and there are many modes, so the broadened resonances overlap and smear into a continuum.

The second problem concerns the significance of all this for biological function. It's very impressive to see quantum tunneling in a biological molecule, but our excitement should be tempered by the fact that we see this only at temperatures below 100 K, far out of the range where life actually happens. Measurements on the (much faster) initial steps of electron transfer, however, show that approximately temperature independent reaction rates persist up to room temperature. Indeed if we look closely at the rates of $P^*I \rightarrow P^+I^-$ and $I^-Q \rightarrow IQ^-$, we see a slightly inverse temperature dependence, with the rate slowing by a factor of two or three as we increase the temperature from 4 to 300 K [should have a figure for this!]. In fact the theory as we have sketched it provides a possible explanation for this: if we tune the energy difference between reactants and products so as to maximize the reaction rate, we have $\epsilon = \lambda$ and the exponential dependence of the reaction rate on T_{eff} disappears; all we have left is $k \propto 1/\sqrt{T_{\text{eff}}}$, which indeed is a weak, inverse temperature dependence. This sort of fine tuning might make sense—perhaps evolution has selected for molecular parameters that maximize the electron transfer rates.

The structure if the reaction center is such that one can take out the quinone molecules and replace them with analogs that have different electron affinities, and in this way manipulate the value of ϵ . Perhaps surprisingly, increases in ϵ have very little effect on the rate constant for the recombination reaction $P^+Q^- \rightarrow PQ$, or on the forward reaction $I^-Q \rightarrow IQ^-$, and for all the values of ϵ probed one sees an approximately temperature independent rate. This argues strongly against tuning of $\epsilon = \lambda$ as an explanation for the observed “activationless” behavior.




FIG. 42 Electron transfer coupled to high frequency vibrations, from Eq's (293) and (294). Dashed lines show contributions to the rate constant at $T = 300\text{ K}$ from processes that leave behind $n = 0, 1, 2, \dots$ quanta in the high frequency mode. The total rate $k(T = 300\text{ K})$ is shown in the solid blue line, and $k(T = 30\text{ K})$ in green. The high frequency mode has $\hbar\Omega = 0.1\text{ eV}$ and $S = 1$.

Suppose that instead of one vibrational mode, we have two—one at a low frequency ω , which we can treat by the semi-classical argument given above, and one at a high frequency Ω that really needs a proper quantum mechanical description. The initial state of the high frequency mode is the ground state (since $k_B T \ll \hbar\Omega$), but in the final state we can excite one or more vibrational quanta, and the overall reaction rate will be a sum over terms corresponding to each of these possible final states. From the point of view of the low frequency mode, if the system transitions into a state with n high frequency quanta, this renormalizes the matrix element $\Delta \rightarrow \Delta_n$ and reduces the energy gap $\epsilon \rightarrow \epsilon - n\hbar\Omega$. Thus the rate constant becomes

$$k = \sum_{n=0}^{\infty} \frac{\Delta_n^2}{\hbar\sqrt{4\pi\lambda k_B T_{\text{eff}}}} \exp\left[-\frac{(\epsilon - n\hbar\Omega - \lambda)^2}{4\lambda k_B T_{\text{eff}}}\right], \quad (293)$$

where now λ refers only to the reorganization energy of the low frequency mode. Results are shown in Fig 42.

Problem 49: Renormalized matrix elements. To complete the calculation in Eq (293), we need to understand how the matrix elements are renormalized by coupling to the high frequency modes. *Get the students to derive ...*

$$\Delta_n^2 = e^{-S} \frac{S^n}{n!} \Delta^2, \quad (294)$$

and explain the meaning of S .

We see that the possibility of exciting different numbers of vibrational quanta greatly broadens the dependence of the rate constant on the energy gap ϵ , and provides a huge widening of the region over which we see very little (or even inverted) temperature dependence. This seems a more plausible and robust explanation of the observed activationless kinetics in the photosynthetic reaction center. Importantly, it relies in an essential way on the quantum behavior of the high frequency vibrational motions that are coupled to the electron transfer, and this is true even at room temperature. There is no shortage of such high frequency modes in the quinones, chlorophylls and pheophytins; what is interesting is the way in which the interplay of these quantum modes with the lower frequency classical modes (including, presumably, modes of the protein scaffolding itself) shapes the observed functional behavior.

A third issue is that, although we are talking about electron transfer reactions, we have said relatively little about the electrons themselves—there are two states, localized on the donor and acceptor sites, and there is a matrix element that connects these states, but that seems to be all. In fact we can say a bit more. First, our use of perturbation theory obviously depends on the matrix element not being too large. If we go back to our simple model of the DeVault–Chance reaction and try to fit the absolute rate constants as well as the temperature dependence, we find $\Delta \sim 10^{-4}\text{ eV}$. Certainly this is small compared with the other energies in the problem ($\lambda, \hbar\omega, k_B T, \epsilon$), which indicates that our use of perturbation theory is consistent. *[Finish the discussion of matrix elements!]*

[Do we want to say anything about coherence and the very first, fastest steps??]

All other things being equal, quantum effects are stronger for lighter particles. As we have seen, electrons essentially always tunnel—there are almost no chemical or biochemical reactions involving thermal activation of an electron over a barrier. Since the early days of quantum mechanics, people have wondered if chemical reactions involving the next lightest particle, a proton or hydrogen atom, might also involve tunneling in a significant way. To be concrete, consider the situation in Fig 43, where the reaction coordinate is the position of the H atom itself, moving from donor to acceptor atom. But, while still attached to the donor atom (e.g., a carbon) we can observe vibrations of the D–H bond, and for C–H we know that the frequencies of these vibrations can be as high as $\nu \sim 2500 - 3000\text{ cm}^{-1}$. The vibrational quanta thus are $h\nu \sim 1/4 - 1/3\text{ eV}$. In fact the activation energies of many chemical reactions are not that much larger




FIG. 43 Transfer of a hydrogen atom from a donor D to an acceptor A. The reaction coordinate is the position of the H atom, but we expect that quantization effects are non-negligible.

than this, perhaps $0.5 - 1$ eV. This means that, as indicated in the crude sketch, climbing up to the top of the barrier between reactants and products involves adding just two or three vibrational quanta. What this means is that the reaction can't really be completely classical, if the reaction coordinate really is the stretching of the bond itself.

If we make the crude approximation that the barrier is rectangular, with height E , then the rate of going over the barrier should be $k \propto e^{-E/k_B T}$, as before, while the

rate of tunneling through the barrier is $k \propto e^{-2\sqrt{2mE}\ell/\hbar}$, where ℓ is the width of the barrier and m is the mass of tunneling particle. Although we could worry about the prefactors, the exponentials are probably the dominant effects, and so we might guess (as in the problem above) that tunneling is more important than classical thermal activation only if

$$e^{-E/k_B T} < e^{-2\sqrt{2mE}\ell/\hbar}, \quad (295)$$

or

$$T < T_0 \sim \frac{\hbar}{k_B} \sqrt{\frac{E}{2m\ell}}. \quad (296)$$

If the width of the barrier is $\ell \sim 1$ Å, and its height is $E \sim 50$ kJ/mole, then with m the mass of the proton we find $T_0 \sim 190$ K, well below room temperature. Thus, although it might be difficult to see the transfer of a proton as being completely classical, it's also true that the transfer reaction is unlikely to be dominated by tunneling at room temperature if the barrier is static.

In the interior of a protein, we can imagine that the donor and acceptor are held by different parts of the large molecule, and as the protein flexes and breathes, these sites will move. Effectively this means that the width of the barrier will fluctuate. On average, this increases the probability of tunneling through the barrier. If the fluctuations in ℓ are Gaussian, the tunneling probability becomes

$$e^{-2\sqrt{2mE}\bar{\ell}/\hbar} \rightarrow \left\langle e^{-2\sqrt{2mE}\ell/\hbar} \right\rangle = \exp \left[-2\sqrt{2mE}\bar{\ell}/\hbar + 4mE\langle(\delta\ell)^2\rangle/\hbar^2 \right], \quad (297)$$

where $\bar{\ell}$ is the average width of the barrier and $\langle(\delta\ell)^2\rangle$ is the variance of this width. With the parameters as before, the enhancement of the tunneling probability involves the term

$$4mE\langle(\delta\ell)^2\rangle/\hbar^2 \sim 6 \left\langle \left(\frac{\delta\ell}{0.1 \text{ Å}} \right)^2 \right\rangle. \quad (298)$$

As described in Appendix A.5, measurements of Debye-Waller factors in X-ray diffraction from protein crystals provide estimates of the fluctuations in structure, and these structural fluctuations are easily several tenths of an Ångström. Thus this term, which appears in the exponential, can be huge. This completely shifts the balance between tunneling and classical, thermal activation, so that in the presence of fluctuations it becomes plausible that tunneling is dominant at room temperature.

Notice that the role of protein vibrational motions here is very different than in the case of electron transfer. In electron transfer, there is a small matrix element that couples the two relevant states, and protein motions serve to bring these two states into degeneracy with one another. This effect presumably could happen in the case of proton transfer as well, but we have focused on the coupling of fluctuations to the tunneling matrix element. This coupling is especially interesting because it generates exponential terms in the reaction rate that have a dependence on mass ($\ln k \propto m$) that is very different from the naive tunneling exponent ($\ln k \propto -\sqrt{m}$) or the zero-point corrections to the activation energy ($\ln k \propto 1/\sqrt{m}$; see next problem); because this mass-dependent term also depends on the variance of structural fluctuations, it is also temperature dependent. Indeed, it was the discov-

ery of anomalous, temperature dependent isotope effects in enzymatic proton transfer reactions that prompted renewed discussion of these dynamical effects on tunneling.

Problem 50: Isotope effects. Chemical reaction rates change when we substitute one isotope for another. There is a “semiclassical” theory of these isotope effects, which says that the reaction proceeds by conventional thermal activation, but the activation energy is reduced by the zero-point energy of vibrations along the reaction coordinate, $k \propto \exp[-(E_{\text{act}} - \hbar\omega)/k_B T]$.

(a.) Vibrational frequencies are proportional to $1/\sqrt{m}$, with m the (effective) mass of the particle(s) moving along the mode with frequency ω . In the simple picture where all of the motion along the reaction coordinate is dominated by the motion of the proton, derive a relationship between the ratios of rate constants for hydrogen, deuterium and tritium transfer.

(b.) If the reaction coordinate involves motion of atoms other than transferred hydrogen, what happens to the predicted magnitude of the isotope effects? What about the relationship you derived in [a.]?

(c.) [Let’s do something with averaging over fluctuating barriers to see how isotope effects come out ...]

I hope that you take a few lessons away from this (long) discussion. First, chemical reactions are the result of fluctuations at the molecular level. We can describe the nature of these fluctuations in some detail, since rare events such as escape over a high barrier are dominated by specific trajectories. In large biological molecules, the flexibility of the molecule means that there is another way for fluctuations to be important, as the variations in protein structure, for example, couple to changes in the barrier for the relevant chemical rearrangements or bring weakly coupled electronic states into degeneracy. Finally, these fluctuations in protein structure can completely revise our view of whether the reaction itself proceeds via classical ‘over the barrier’ motion or by quantum tunneling. These theoretical observations, and the experiments to which they connect, suggest that Nature exploits not just the structure of biological molecules, but also the fluctuations in these structures, to control the rates of chemical reactions.

If you need a review of the Langevin equation, I like the treatment in the little book by Kittel (1958), as well as the somewhat longer discussion by Pathria (1972). Every physics student should understand the basic instanton calculation of tunneling, as an illustration of the power of path integrals. There is no better treatment than that given by Coleman in his justly famous Erice lectures. If you read Coleman you’ll not only get a deeper view of what we have covered here, you’ll get all the missing pieces about the prefactor of the rate constant, and much more. For more general background on path integrals, including some discussion of how to use them for classical stochastic processes, the standard reference is Feynman &

Hibbs (1965). For more rigorous accounts of many of these issues (e.g., getting the Jacobian right in constructing the path integral), see Zinn–Justin (1989). The original discussion of diffusion (even with inertia) over a barrier is due to Kramers (1940); for a modern perspective see Hänggi et al (1990).

Coleman 1988: *Aspects of Symmetry* S Coleman (Cambridge University Press, Cambridge, 1988).

Feynman & Hibbs 1965: *Quantum Mechanics and Path Integrals* RP Feynman & AR Hibbs (McGraw–Hill, New York, 1965).

Hänggi et al 1990: Reaction–rate theory: Fifty years after Kramers. P Hänggi, P Talkner & M Borkovec, *Revs Mod Phys* **62**, 251–341 (1990).

Kittel 1958: *Elementary Statistical Physics* C Kittel (Wiley, New York, 1958).

Kramers 1940: Brownian motion in a field of force and the diffusion model of chemical reactions. HA Kramers *Physica* **7**, 284–304 (1940).

Pathria 1972: *Statistical Mechanics* RK Pathria (Pergamon Press, Oxford, 1972).

Zinn–Justin 1989: *Quantum Field Theory and Critical Phenomena* J Zinn–Justin (Oxford University Press, Oxford, 1989).

Myoglobin was the first protein whose structure was solved by X–ray diffraction. Aspects of X–ray analysis are described in Appendix A.5. For a perspective on myoglobin, see Kendrew (1964). The experiments on myoglobin are by Austin et al (1975), which touched off a huge followup literature. A clear discussion of the interplay between the a reaction coordinate and a protein coordinate was given by Agmon and Hopfield (1983). The demonstration of tunneling in this system is by Alberding et al (1976).

Agmon & Hopfield 1983: Transient kinetics of chemical reactions with bounded diffusion perpendicular to the reaction coordinate. N Agmon & JJ Hopfield, *J Chem Phys* **78**, 6947–6959 (1983).

Alberding et al 1976: Tunneling in ligand binding to heme proteins. N Alberding, RH Austin, KW Beeson, SS Chan, L Eisenstein, H Frauenfelder & TM Nordlund, *Science* **192**, 1002–1004 (1976).

Austin et al 1975: Dynamics of ligand binding to myoglobin. RH Austin, KW Beeson, L Eisenstein, H Frauenfelder & IC Gunsalus, *Biochemistry* **14**, 5355–5373 (1975).

Kendrew 1964: Myoglobin and the structure of proteins. JC Kendrew, in *Nobel Lectures in Chemistry 1942–1962* (Elsevier, Amsterdam, 1964). See also <http://www.nobelprize.org>.

Classical overviews of the photosynthetic reaction center are provided by Feher & Okamura (1978) and Okamura et al (1982). As with many biological molecules, many questions about the reaction center were sharpened once the structure was determined at atomic resolution (Deisenhoffer et al 1984); this work was important also as a demonstration that one could use the classical methods of X–ray crystallography (cf Appendix A.5) for proteins that are normally embedded in membranes. It should be emphasized, however, that the electron transfer reactions leave an enormous variety of spectroscopic signatures—separating charges not only changes optical properties of the molecules, it generates unpaired spins that can be seen using electron paramagnetic resonance (EPR), and the distribution of the spin across multiple atoms at the donor and acceptor sites can be mapped using electron–nuclear double resonance (ENDOR). An early view of the uses of EPR and ENDOR in biological systems is given by Feher (1970); this article appears in the proceedings of the first Les Houches physics summer school to be devoted to questions at the interface with biology [check this!]. For a synthesis of structural and spectroscopic data in relation to function, see Feher et al (1989).

- DeVault & Chance 1966:** Studies of photosynthesis using a pulsed laser. I. Temperature dependence of cytochrome oxidation in *Chromatium*. Evidence for tunneling. D DeVault & B Chance, *Biophys J* **6**, 825–847 (1966).
- Deisenhoffer et al 1984:** X-ray structure analysis of a membrane protein complex: Electron density map at 3 Å resolution and a model of the chromophores of the photosynthetic reaction center from *Rhodospseudomonas viridis*. J Deisenhoffer, O Epp, K Miki, R Huber & H Michel, *J Mol Biol* **180**, 385–398 (1984).
- Feher 1970:** Electron paramagnetic resonance with applications to selected problems in biology. G Feher, in *Physical Problems in Biological Systems*, C DeWitt & J Matricon, eds, pp 251–365 (Gordon & Breach, Paris, 1970).
- Feher & Okamura 1978:** Chemical composition and properties of reaction centers. G Feher & MY Okamura, in *The Photosynthetic Bacteria*, RK Clayton & WR Sistrom, eds, pp 349–386 (Plenum Press, New York, 1978).
- Feher et al 1989:** Primary processes in bacterial photosynthesis: structure and function of reaction centers. G Feher, JP Allen, MY Okamura & DC Rees, *Nature* **339**, 111–116 (1989).
- Okamura et al 1982:** Reaction centers. MY Okamura, G Feher & N Nelson, in *Photosynthesis: Energy Conversion by Plants and Bacteria, Volume 1*, Govindjee, ed, pp 195–272 (Academic Press, New York, 1982).

The original experiments that provided evidence for tunneling in photosynthetic electron transfer were done by DeVault and Chance (1966) on samples that were a bit messier than the purified reaction centers that emerged in subsequent years. The kinetics of the initial charge separation reactions were described by [fill in refs!]. The modern view of biological electron transfer reactions, including the role of tunneling in the vibrational degrees of freedom, is due to Hopfield (1974). Exploration of the energy gap dependence of reaction rates was pioneered by Gunner et al (1986), and the evidence for frozen distributions of electron transfer rates was provided by Kleinfeld et al (1984). For a review of efforts to calculate electronic matrix elements in real protein structures, see Onuchic et al (1992). [Maybe we need more here? Depends also on what happens in the text.]

- DeVault & Chance 1966:** Studies of photosynthesis using a pulsed laser. I. Temperature dependence of cytochrome oxidation in *Chromatium*. Evidence for tunneling. D DeVault & B Chance, *Biophys J* **6**, 825–847 (1966).
- Gunner et al 1986:** Kinetic studies on the reaction center protein from *Rhodospseudomonas sphaeroides*: the temperature and free energy dependence of electron transfer between various quinones in the Q_A site and the oxidized bacteriochlorophyll dimer. MR Gunner, DE Robertson & PL Dutton, *J Phys Chem* **90**, 3783–3795 (1986).
- Hopfield 1974:** Electron transfer between biological molecules by thermally activated tunneling. JJ Hopfield, *Proc Nat'l Acad Sci (USA)* **71**, 3640–3644 (1974).
- Kleinfeld et al 1984:** Electron-transfer kinetics in photosynthetic reaction centers cooled to cryogenic temperatures in charge-separated state: evidence for light-induced structural changes. D Kleinfeld, MY Okamura & G Feher, *Biochemistry* **23**, 5780–5786 (1984).
- Onuchic et al 1992:** Pathway analysis of protein electron-transfer reactions. JN Onuchic, DN Beratan, JR Winkler & HB Gray, *Annu Rev Biophys Biomol Struct* **21**, 349–377 (1992).

The papers that reignited interest in proton tunneling in enzymes were Cha et al (1989) and Grant & Klinman (1989). The idea that these experiments should be understood in terms of coupling

between quantum motion of the proton and classical motion of the protein was developed by Bruno & Bialek (1992). It took roughly a decade for these ideas to solidify, as described in reviews by Sutcliffe & Scrutton (2002) and Knapp & Klinman (2002). [add refs to Nori et al in cytochrome oxidase?]

- Bruno & Bialek 1992:** Vibrationally enhanced tunneling as a mechanism for enzymatic hydrogen transfer. WJ Bruno & W Bialek, *Biophys J* **63**, 689–699 (1992).
- Cha et al 1989:** Hydrogen tunneling in enzyme reactions. Y Cha, CJ Murray & JP Klinman, *Science* **243**, 1325–1330 (1989).
- Grant & Klinman 1989:** Evidence that both protium and deuterium undergo significant tunneling in the reaction catalyzed by bovine serum amine oxidase. KL Grant & JP Klinman, *Biochemistry* **28**, 6597–6695 (1989).
- Knapp & Klinman 2002:** Environmentally coupled hydrogen tunneling: Linking catalysis to dynamics. MJ Knapp & JP Klinman, *Eur J Biochem* **269**, 3113–3121 (2002).
- Sutcliffe & Scrutton 2002:** A new conceptual framework for enzyme catalysis: Hydrogen tunneling coupled to enzyme dynamics in flavoprotein and quinioprotein enzymes. MJ Sutcliffe & NS Scrutton, *Eur J Biochem* **269**, 3096–3102 (2002).

B. Molecule counting

Many of the crucial signals in biological systems—signals that are internal to cells, signals that cells use to communicate with one another, even signals that organisms exchange—are carried by changes in the concentration of specific molecules. The molecules range in size from single ions (e.g., calcium) to whole proteins. Such chemical signals act by binding to specific targets, whose synthesis and accessibility can also be controlled by the cell. A key point is that individual molecules move randomly, and so the arrival of signals at their targets has some minimum level of noise. As we shall see, several different systems operate with a reliability close to this physical limit: in essence, these systems are counting every molecule, and making every molecule count.

In what follows we will see examples of chemical signaling in the decisions that cells make about whether to read out the information encoded in particular genes, in the trajectories that axons take toward their targets in the developing brain, in the control signals that bacteria use to regulate their movement, and in the development of spatial patterns in a developing embryo. But much of our thinking about precision, reliability and noise in chemical signaling has been shaped by the phenomena of chemotaxis in bacteria, so this is where we will start.

Although our experience with other animals makes it clear that we are not alone in our ability to sense the world, it still seems remarkable that single celled organisms such as bacteria are endowed with sensory systems

that allow them to move in response to a variety of signals from the environment, including the concentrations of various chemicals. A classical observation (from the 19th century) is that some bacteria, swimming in water on a microscope slide, under a cover slip, will collect at the center of cover slip, while others will collect at the edges. Those with more refined tastes will form a tight band that traces the outlines of the square cover slip. Oxygen diffuses into the water through the edges of the cover slip, and by collecting along a square the bacteria have migrated to a place of constant (not maximal or minimal) oxygen concentration. It is plausible that this happens because they can sense the oxygen concentration and “know” the most comfortable value of this concentration, much as we might move to be the most comfortable distance from a fireplace in an otherwise unheated room.

That bacteria collect at nontrivial concentrations of different molecules really doesn't demonstrate that they sense the concentration. They might instead sense some internal consequences of the external variables, such as the accumulation of metabolic intermediates. In the 1960s Adler found mutants of *E coli* which cannot metabolize certain sugars or amino acids but will nevertheless migrate toward the sources of these molecules; also there are mutants that metabolize but can't migrate. This is convincing evidence that metabolism and sensing are separate systems, and thus begins the fruitful exploration of the sensory mechanisms of bacteria and the connection of these sensory mechanisms to motor output. This phenomenon is called *chemotaxis*.

I'll skip lots of the truly classical stuff and proceed with the modern biophysical approach, which begins ~ 1970 . To a large extent this modern approach rests on the work of Howard Berg and collaborators. The first key step taken by Berg and Brown was to observe the behavior of individual bacteria. *E coli* are $\sim 1\ \mu\text{m}$ in size, and can be seen relatively easily under the light microscope, but since the bacteria swim at ~ 20 body lengths per second they easily leave the field of view or the plane of focus; the solution is to build a tracking microscope.

Observations in the tracking microscope, as in Fig 44, showed that the trajectories of individual bacteria consist of relatively straight segments interrupted by short intervals of erratic “motion in place.” These have come to be called runs and tumbles, respectively. Tumbles last ~ 0.1 seconds, but the erratic motion during this brief time is sufficient to cause successive runs to be in almost random relative directions. Thus the bacterium runs in one direction, then tumbles and chooses a new direction at random, and so on. Runs themselves are distributed in length, as if the termination of a run is itself a random process.

Closer examination of the runs shows how it is possible for this seemingly random motion to generate progress up the gradient of attractive chemicals. When the bac-




FIG. 44 Paths of *E coli* as seen in the original tracking microscope experiments, from Berg & Brown (1972). The three panels in each case are projections of the path onto the three orthogonal planes (imagine folding the paper into a cube along the dashed lines). At left, wild type bacteria, showing the characteristic runs and tumbles. At right, a non-chemotactic mutant that never manages to tumble.

terium runs up the gradient, the mean duration of the runs becomes longer, biasing the otherwise random walk. Interestingly, when bacteria swim down the gradient (of an attractant, or up the gradient of a repellent) the is relatively little change in the mean run length. Berg has described this as a form of optimism: If things are getting better, keep going, but if things are getting worse, don't worry. [Need to look at the notion of optimism once more in relation to all the data.]

Since runs get longer when bacteria swim along a positive gradient, it is natural to ask whether the cell is responding to the spatial gradient itself or to the change in concentration with time along the path. As we will see, the spatial gradients to which the cell can respond are very small, and searching for a systematic difference (for example) between the front and back of the bacterium is unlikely to be effective just on physical grounds, independent of biological mechanisms. Indeed, this is the reason why chemotaxis is such an important example of the issues in this section. To search for a time domain mechanism one can expose the bacteria to concentrations which are spatially uniform but varying in time; if the sign of the change corresponds to swimming up a positive gradient, runs should be prolonged. The first such experiment used very large, sudden changes in concentration, and found that cells were trapped in extremely long runs. A more sophisticated experiment used enzymes to synthesize attractants from inert precursors, exposing the cells to gradual changes more typical of those encountered while swimming. Purely time domain stimuli were sufficient to generate modulations of run length that agree quantitatively with those observed for bacteria experiencing spatial gradients.

Problem 51: Chemotaxis in one dimension. To make the intuition of the previous paragraphs more rigorous, consider a simplified problem of chemotaxis in one dimension. There are then two populations of bacteria, the + population that moves to the right and the - population that moves to the left, each at speed v . Let the probability of finding a + [-] bacterium at position x

be $P_+(x, t)$ [$P_-(x, t)$]. Assume that the rate of tumbling depends on the time derivative of the concentration along the bacterial trajectory as some function $r(\dot{c})$, where for the \pm bacteria we have $\dot{c} = \pm v dc/dx$, and that cells emerge from a tumble going randomly left or right.

(a.) Show that the dynamics of the two probabilities obey

$$\frac{\partial P_+(x, t)}{\partial t} + v \frac{\partial P_+(x, t)}{\partial x} = -r \left(v \frac{dc}{dx} \right) P_+(x, t) + \frac{1}{2} \left[r \left(v \frac{dc}{dx} \right) P_+(x, t) + r \left(-v \frac{dc}{dx} \right) P_-(x, t) \right] \quad (299)$$

$$\frac{\partial P_-(x, t)}{\partial t} - v \frac{\partial P_-(x, t)}{\partial x} = -r \left(-v \frac{dc}{dx} \right) P_-(x, t) + \frac{1}{2} \left[r \left(v \frac{dc}{dx} \right) P_+(x, t) + r \left(-v \frac{dc}{dx} \right) P_-(x, t) \right]. \quad (300)$$

Explain the meaning of each of the terms in terms of what happens as cells enter into and emerge from tumbles. Note that in this approximation tumbles themselves are instantaneous, which isn't so bad (0.1 s vs the $\sim 1 - 10$ s for typical runs).

(b.) To see if the bacteria really migrate toward high concentrations, look for the steady state of these equations. If we simplify and assume that the rate of tumbling is modulated linearly by the time derivative of the concentration,

$$r(\dot{c}) \approx r(0) + \frac{\partial r}{\partial \dot{c}} \dot{c} + \dots, \quad (301)$$

show that

$$P(x) = \frac{1}{Z} \exp \left[-\frac{\partial r}{\partial \dot{c}} c(x) \right]. \quad (302)$$

Thus, in these approximations, chemotaxis leads to a Boltzmann distribution of bacteria, in which the concentration acts as a potential. If the molecules are attractive then $\partial r/\partial \dot{c} < 0$ and hence maxima of concentration are minima of the potential, conversely for repellents. The stronger the modulation of the tumbling rate (as long as we stay in our linear approximation) the lower the effective temperature and the tighter the concentration of bacteria around the local maxima of concentration.

Problem 52: Nonlinearities. Within this simplified one dimensional world, can you make progress without the approximation that $r(\dot{c})$ is linear? More specifically, what is the form of the stationary distribution $P(x)$ that solves Eq (??) for nonlinear $r(\dot{c})$? Can you show that there still is an effective potential with minima located at places where the concentration is maximal?

Problem 53: A little more about the effectiveness of chemotaxis.

(a.) Within the one dimensional model, what happens if the tumbling rate is modulated not just by the time derivative, but also by the absolute concentration, so that the bacterium confuses "currently good" for "getting better"?

(b.) Can you generalize this discussion to three dimensions? Instead of having just two groups + and -, one now needs a continuous distribution $P(\Omega, x, t)$, where Ω denotes the direction of swimming. Derive an equation for the dynamics of $P(\Omega, x, t)$ in the same approximations used above, and see if the Boltzmann-like solution obtains in this more realistic case.

order of one or two meters, roughly the length of our bodies. In contrast, if a bacterium stops running its motors, it will glide for a distance comparable not to its body length ($\sim 1 \mu\text{m}$) but to the diameter of an atom. To see this, think about a small particle moving through a fluid, subject only to drag forces (the motors are off). If the velocities are small, we know the drag will be proportional to the velocity, so Newton's equation is just

$$m \frac{dv}{dt} = -\gamma v. \quad (303)$$

For a spherical object of radius r , Stokes' law tells us that $\gamma = 6\pi\eta r$, where η is the viscosity of the fluid, and we also know that $m = 4\pi\rho r^3/3$, where ρ is the density of the object. The result is that

$$v(t) = v(0) \exp(-t/\tau), \quad (304)$$

where

$$\tau = \frac{m}{\gamma} = \frac{2\rho r^2}{9\eta}. \quad (305)$$

If we assume that the density of bacteria is roughly that of water, then it is useful to recall that η/ρ has units of a diffusion constant, and for water $\eta/\rho = 0.01 \text{ cm}^2/\text{s}$. With $r \sim 1 \mu\text{m} = 10^{-4} \text{ cm}$, this gives $\tau \sim 5 \times 10^{-7} \text{ s}$. If the initial velocity is $v(0) \sim 20 \mu\text{m}/\text{s}$, the net displacement during this coasting is $\Delta x = v(0)\tau \sim 10^{-11} \text{ m}$; recall that a hydrogen atom has a diameter of $\sim 1 \text{ \AA} = 10^{-10} \text{ m}$.

The conclusion from such simple estimates is that bacteria can't coast. More generally, mechanics on the scale of bacteria is such that inertia is negligible, as if Aristotle (rather than Galileo and Newton) were right. This is really about the nature of fluid flow on this scale.³⁶ For an incompressible fluid (which is a good approximation here—surely the bacteria don't generate sound waves as

All of this description so far is about the phenomenology of swimming. But how does it actually work? The basic problem is that bacteria are too small to take advantage of inertia. When we swim, we can push off the wall of the pool and glide for some distance, even without moving our arms or legs; this gliding distance is on the

³⁶ My experience is that most physics students don't know too much fluid mechanics, so although this is elementary I put it here. For a more thorough discussion, see, as usual, Landau and Lifshitz.

they swim), the Navier–Stokes equations are

$$\rho \left[\frac{\partial \mathbf{v}}{\partial t} + \mathbf{v} \cdot \nabla \mathbf{v} \right] = -\nabla p + \eta \nabla^2 \mathbf{v}, \quad (306)$$

where \mathbf{v} is the local velocity of the fluid, p is the pressure, and as usual ρ is the density and η is the viscosity. The pressure isn't really an independent variable, but needs to be there so we can enforce the condition of incompressibility,

$$\nabla \cdot \mathbf{v} = 0. \quad (307)$$

These equations need to be supplemented by boundary conditions, in particular that the fluid moves with the same velocity as any object at the points where it touches that object. Thus the velocity should be zero at a stationary wall, and should be equal to the velocity of a swimmer at the swimmer's surface.

Problem 54: Understanding Navier–Stokes. This isn't a fluid mechanics course, but you should be sure you understand what Eq (306) is saying. In particular, this is nothing but Newton's $F = ma$. Explain.

Dimensional analysis is an enormously powerful tool in fluid mechanics. We are free to choose new units for length (ℓ) and time (t_0), and hence for velocity ($v_0 = \ell/t_0$), as well as for pressure p_0 , and this gives us

$$\rho \left[\frac{v_0}{t_0} \frac{\partial \tilde{\mathbf{v}}}{\partial \tilde{t}} + \frac{v_0^2}{\ell} \tilde{\mathbf{v}} \cdot \tilde{\nabla} \tilde{\mathbf{v}} \right] = -\frac{p_0}{\ell} \tilde{\nabla} \tilde{p} + \eta \frac{v_0}{\ell^2} \tilde{\nabla}^2 \tilde{\mathbf{v}}, \quad (308)$$

$$\frac{\rho \ell v_0}{\eta} \left[\frac{\partial \tilde{\mathbf{v}}}{\partial \tilde{t}} + \tilde{\mathbf{v}} \cdot \tilde{\nabla} \tilde{\mathbf{v}} \right] = -\frac{p_0 \ell}{\eta v_0} \tilde{\nabla} \tilde{p} + \tilde{\nabla}^2 \tilde{\mathbf{v}}, \quad (309)$$

where $\tilde{t} = t/t_0$, $\tilde{\mathbf{v}} = \mathbf{v}/v_0$, and $\tilde{p} = p/p_0$. Now we can set $p_0 \ell / \eta v_0 = 1$, which gets rid of all the units, except we are left with a dimensionless combination

$$\text{Re} \equiv \frac{\rho \ell v_0}{\eta} \quad (310)$$

which is called the Reynolds' number. Notice that if we choose the unit of length to be the size of the objects that we are interested in, and v_0 to be the speed at which they are moving, then even the boundary conditions don't have any units, nor do they introduce any dimensionless factors that are far from unity. The conclusion is that all fluid mechanics problems with the same geometry (shapes) are the same if they have they have

the same Reynolds' number. In this sense, being smaller (reducing ℓ) is the same as living at increased viscosity.³⁷

To make a long story short, we live at high Reynolds' number, and bacteria live at low Reynolds' number (Fig 45). Turbulence is a high Reynolds' number phenomenon, as is the more mundane gliding through the pool after we push off the wall. At low Reynolds' number, life is very different. Inertia is absent, and so forces must balance at every instant of time. To say this more startlingly, if $\text{Re} \rightarrow 0$ then time doesn't actually appear in the equations. This means that, as you swim, the distance that you move depends on the *sequence* of motions that you go through, but not on the dynamics with which you execute them.

To use Purcell's evocative example, at high Reynolds' number a scallop can propel itself by snapping shut, expelling a jet of water, and then opening slowly.³⁸ The jet will propel the scallop forward, and the drag of reopening can be made small by moving slowly. At low Reynolds' number this doesn't work, and the forward displacement generated by snapping shut will be exactly compensated by the drag on reopening. To have net movement from a cycle, the *sequence* of shapes that the swimmer goes through in the cycle must break time reversal invariance,




FIG. 45 Purcell's delightful sketch, illustrating the range of Reynolds' numbers relevant for swimming in humans, fish, and bacteria. From Purcell (1977).

³⁷ It is worth reflecting on the level of universality that we have here. We could imagine starting with a molecular description of fluids, then figuring out that, on the relevant length and time scales, all we need to know are the density and viscosity. Now we see that even these quantities are tied up with our choice of units. If we want to know what happens in natural units (i.e., scaling to the size and speed of the objects we are looking at), then all that matters is a single dimensionless combination, Re .

³⁸ There is an interesting issue about what real scallops do. Check Rob's note about this!

not just the trajectory.

So, how do bacteria evade the “scallop theorem”? If you watch them swimming, you can see that the have long filaments sticking out, and these seem to be waving. I emphasize that “see” is tough here. [This needs pictures; check with Berg.] These filaments are very small, ~ 20 nm in diameter, much thinner than the wavelength of light. To see them, the easiest thing is to use dark field microscopy, in which the sample is illuminated from the side and what you see is the light scattered by $\sim 90^\circ$. These apparently waving appendages are called flagella, and remind us of what we see on other small swimming cells, such as sperm. The difference is that the flagella in these other cases are huge by comparison with the bacterial flagella. If you slice through the tail of a sperm and take an electron micrograph, you find an enormously complex structure, and if you try to analyze the system biochemically you find it is made from many different proteins. Importantly, some of these proteins act as enzymes and eat ATP, which we know is a source of energy, for example in our muscles. In contrast, the bacterial flagellum is small, with a relatively simple structure, and the biochemistry suggests that it is little more than a very long polymer made from one kind of protein; this protein is not an enzyme. How can this simple structure, with no ATPase activity, generate motions?

In experiments that aimed at better ways to see the flagella, one can attach “flags” to them using viruses that would stick to the flagella via antibodies. Once in a while, a virus with antibodies on both ends would stick to two flagella from different bacteria. When this happened, you could see the bacterial cells rotating, which one can imagine was a huge surprise. Eventually people figured out how to break off the flagella and stick the bacteria to a glass slide by the remaining stump, and then the bacterium rotates. Rotation can look like a wave if the flagellum is shaped like a corkscrew, and it is. Rotating a corkscrew obviously violates time reversal invariance. If you have several corkscrews and you rotate them with the correct handedness, they can fit together into a bundle. If you rotate the other way, the corkscrews clash, and any bundle will be blown apart by this clashing. So, with many flagella projecting from their surface, we can imagine that by switching the direction of rotation, the bacterium switches between a bundle that can smoothly propel the cell forward, and many independently moving flagella that would cause the cell to tumble in place—runs and tumbles correspond to counterclockwise and clockwise flagellar rotation.³⁹ If you find mutants that never tumble, and stick them down by their stumps, then they

all rotate one way; similarly, mutants that tumble too often rotate the other way.

There is much more to say about the rotary engine itself, sitting at the base of the flagella. It is powered not by ATP but by a difference in chemical potential for hydrogen ions between the inside and the outside of the cell. This is an energy source that all cells use, albeit in different ways, because it allows chemical events at very different spatial locations to be coupled. Thus, as described in the preceding section, photosynthetic organisms use the energy of the absorbed photons to move electrons across a membrane, and then compensate the charges by moving protons; the resulting difference in chemical potential can be used by other membrane-spanning enzymes to make ATP, without being anywhere near the molecules that absorb the photon.⁴⁰ In fact, these enzymes that synthesize ATP also rotate as they let protons move down the gradient in their chemical potential, and these same enzymes are responsible for ATP synthesis in all cells. So, proton driven rotary motors are at the heart of energy conversion in all organisms.

There is also more to say about mechanics at low Reynolds’ number. Swimming involves changing shape, and this provides the boundary conditions on the Navier–Stokes equations. A cycle of changing boundary conditions should lead to a net displacement. There is some subtlety here, since the space of shapes is not so easy to parameterize. If we think, for example, about a closed surface, “shape” is defined by three dimensional position as a function of the two coordinates on the surface (e.g., latitude and longitude), but there is an arbitrariness in how we choose these coordinates; of course any physical quantity, such as the amount by which the swimmer moves forward, must be invariant to this choice. Looking more closely, the freedom to choose coordinates means that the natural formulation of the problem includes a gauge symmetry. Reluctantly, let’s leave all this and go back to the problem of chemotaxis itself.

Problem 55: Switching in tethered bacteria. As noted above, one way of studying bacterial motility and chemotaxis is to “tether” a bacterium by the stump of one flagellum, observing the rotation of the whole cell rather than the rotation of the flagellum. The file `omega.txt` contains a very long time series of the angular velocity from such an experiment done by WS Ryu, now at the University of Toronto.⁴¹ The samples are taken sixty times per second, and the units of velocity are not quite arbitrary but

³⁹ This association goes of course depends on our convention for defining the handedness of rotation; it doesn’t matter (and I have trouble remembering it!) as long as you are consistent.

⁴⁰ You can imagine how confusing this was before people figured it out! It looked like a mysterious action at a distance.

⁴¹ Data that you need can be found at <http://www.princeton.edu/~wbialek/PHY562/data.html>.

[What is the permanent way of dealing with this??]

not really important either; you should be able to load this into MATLAB (load omega.txt).

(a.) You should see that the velocity switches between positive and negative values, but these values are fairly constant. This is consistent with swimming by switching between runs and tumbles, with little or no modulation of the swimming speed. What is the distribution of times spent with during each segment of positive or negative (clockwise or counterclockwise) velocity?

(b.) It usually is said that switching is a Poisson process, so that (as you remember from the discussion of photon counting) the distribution of intervals between switches should be exponential. Are your results in [a] consistent with this prediction?

(c.) Look carefully at the velocity vs. time in the data set. Are the data statistically stationary (time–translation invariant)? If you focus on segments of the data that are more clearly stationary, does that change your conclusions in [b]?

(d.) Sometimes the angular velocity makes a “partial switch,” a brief excursion away from the typical positive or negative value but not quite a full switch to the opposite direction of rotation. Qualitatively, what is happening in these cases? What would be the simplest model to describe the velocity vs. time during such an event? Can you give a quantitative analysis of the data, fitting to your model? This is a bit open ended.

We are interested in the question of how sensitively the bacterium can respond to small concentration gradients. We suspect that, since individual molecular motions are random, there must be a limit, analogous to the shot noise in counting photons. In a classic paper, Berg and Purcell provided a clear intuitive picture of the noise in ‘measuring’ chemical concentrations. Their argument, schematized in Fig 46, was that if we have a sensor with linear dimensions a , then effectively the sensor samples a volume a^3 . In this volume we expect to count an average of $N \sim ca^3$ molecules when the concentration is c . Each




FIG. 46 A schematic of concentration measurements. A receptor of linear dimension a samples a volume a^3 and hence sees a mean number of molecules $N = ca^3$, where c is the concentration. These molecules random walk in and out of the sensitive volume with a diffusion constant D , corresponding to an equilibration or correlation time $\tau_c = a^2/D$.

such measurement, however, is associated with a noise $\delta N_1 \sim \sqrt{N}$. Since the count of molecules is proportional to our estimate of the concentration, the fractional error will be the same, so from one observation we obtain a precision

$$\frac{\delta c}{c} \Big|_1 = \frac{\delta N_1}{N} = \frac{1}{\sqrt{N}} = \frac{1}{\sqrt{ca^3}}. \quad (311)$$

We can make more accurate measurements by averaging over time, although this is a bit tricky—we won’t get a better estimate of the concentration around us by counting the same molecules over and over again. Thus if we are willing to average over a time τ_{avg} , we can make K independent measurements, where $K \sim \tau_{\text{avg}}/\tau_c$, and the correlation time τ_c is the time we have to wait in order to get an independent sample of molecules.

How do we get independent samples? If we look in a small volume, the molecules that we are looking at exchange with the surroundings through diffusion. Thus the time required to get an independent collection of molecules is the time required for molecules to diffuse in and out of the volume, $\tau_c \sim a^2/D$. Putting everything together we have

$$\frac{\delta c}{c} = \frac{1}{\sqrt{K}} \cdot \frac{\delta c}{c} \Big|_1 \quad (312)$$

$$= \sqrt{\frac{\tau_c}{\tau_{\text{avg}}}} \cdot \frac{1}{\sqrt{ca^3}} \quad (313)$$

$$= \sqrt{\frac{a^2}{D\tau_{\text{avg}}}} \cdot \frac{1}{\sqrt{ca^3}} \quad (314)$$

$$= \frac{1}{\sqrt{Dac\tau_{\text{avg}}}}. \quad (315)$$

This is a lovely result. It says that the limit to the accuracy of measurements depends on the absolute concentration (more molecules \rightarrow more accuracy), on the size the detector (bigger detectors \rightarrow more accuracy), on the time over which we are willing to average (more time \rightarrow more accuracy), and finally on the diffusion constant of the molecules we are sensing, because faster diffusion lets us see more independent samples in the same amount of time. All these parameters combine simply, essentially in the only way allowed by dimensional analysis.

One way of understanding this result on the limits to precision is to think about the rate at which molecules find their target. For molecules at concentration c moving with diffusion constant D , the rate (number of molecules per second) that arrive at a target of size a should be proportional both to c and to D , and then by dimensional analysis we need one factor of length, so the rate is $\sim Dac$ molecules per second. This result is used most often to talk about the “diffusion limited rate constant” for a chemical reaction; if we have



then the second order rate constant k_+ can never be bigger than $\sim Da$, where D is the diffusion constant of the molecules and a is their size, or more precisely the size of the region where they have to hit in order to react. But if the rate of molecular arrivals is $\sim Dac$, in a time τ_{avg} we will count $\sim Dac\tau_{\text{avg}}$ molecules, and if these molecules are arriving at random then there will be the usual square root fluctuations, which leads us to Eq (315). In this view, the Berg–Purcell limit is nothing but shot noise in molecular arrivals, and thus is completely analogous to shot noise in photon arrivals. Photons propagate and molecules diffuse, but under most conditions they both arrive at random, hence there is shot noise in counting.

Problem 56: Diffusion limited rates, more carefully. One can try a more careful calculation of the rate at which molecules find their target by diffusion. Imagine a sphere of radius a such that all molecules which hit the surface are immediately absorbed. Outside the sphere, the concentration profile must obey the diffusion equation, and the absorption means that on the spherical surface the concentration will be zero. Far from the sphere, the concentration should be equal to c . Thus we have

$$\frac{\partial c(\mathbf{x}, t)}{\partial t} = D\nabla^2 c(\mathbf{x}, t); \quad (317)$$

$$c(|\mathbf{x}| = a, t) = 0, \quad (318)$$

$$c(\mathbf{x} \rightarrow \infty, t) = c. \quad (319)$$

The number of molecules arriving per second at the surface of the sphere is given by an integral of the diffusive flux over the surface

$$\text{rate} = \int d^2s \hat{\mathbf{n}} \cdot [-D\nabla c(\mathbf{x}, t)] \Big|_{|\mathbf{x}|=a}, \quad (320)$$

where d^2s is an element of the surface area on the sphere, and $\hat{\mathbf{n}}$ is the unit vector normal to the sphere.

(a.) Solve Eq (317), with the boundary conditions in Eqs (318) & (319), in steady state. Note that as a first step you should go to spherical coordinates; recall that in three dimensions the Laplacian can be written as

$$\nabla^2 = \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial}{\partial r} \right) + \frac{1}{r^2} \left[\frac{1}{\sin^2 \phi} \frac{\partial^2}{\partial \theta^2} + \frac{1}{\sin \phi} \frac{\partial}{\partial \phi} \left(\sin \phi \frac{\partial}{\partial \phi} \right) \right], \quad (321)$$

where as usual r is the radius and θ and ϕ are the polar and azimuthal angles, respectively.

(b.) Use your steady state solution to evaluate the rate at which molecules arrive at the sphere, using Eq (320). Also, explain why simple dimensional analysis of these equations yields rate $\sim Dac$.

(c.) What happens if you try to give a dimensional analysis argument for the rate in one or two dimensions? If there are problems, can you explain how these problems either go away or are made more precise by trying to solve the diffusion equation with appropriate boundary conditions? As a hint, the two dimensional case is a bit delicate; focus first on one dimension.

Bacteria such as *E. coli* have been observed to perform chemotaxis in environments where ambient concentrations of attractants such as sugars or amino acids are as low as ~ 1 nM, which is $\sim 10^{-9} \times (6 \times 10^{23})/10^3 =$

6×10^{11} molecules per cm^3 . These small molecules diffuse through aqueous solution with $D \sim 10^{-5} \text{cm}^2/\text{s}$, and the most generous assumption would be that the relevant size of the detector is the size of the whole bacterium, $a \sim 1 \mu\text{m}$. Putting these factors together, we have $Dac \sim 600 \text{s}^{-1}$. Thus, if the bacterium integrates for $\tau_{\text{avg}} \sim 1.5 \text{s}$, the smallest concentration changes it can detect are $\delta c/c \sim 1/30$. If the cells were to detect the difference in concentrations across the $\sim 1 \mu\text{m}$ length of their body, this would mean that the concentration was varying significantly on the scale of $30 \mu\text{m}$, which is very short indeed. In real experiments (and, presumably, in the natural environment) the length scales of concentration gradients are one to two orders of magnitude longer. Thus, it's impossible—without integrating for minutes or hours—for bacteria to perform as they do by measuring a spatial gradient. The only possibility is to measure the concentration variation in time, along the trajectory that the bacterium takes through the gradient. Since the cells move at $v = 10 - 20 \mu\text{m}/\text{s}$, on times scales of $\tau_{\text{avg}} \sim 1.5 \text{s}$ this increases the signal by a factor of ten to thirty, and brings the signal above the background of noise, allowing for reliable detection.

[Maybe add remarks that this argument still works at higher concentrations, if the length scales of gradients are even longer? Perhaps this could be put into a problem?]

Although the comparisons are a bit rough,⁴² we can draw several conclusions. First, real bacteria perform chemotaxis in response to small signals with a reliability close to the limits set by the physics of diffusion. Second, this is possible only if the cell measures the derivative of concentration vs. time as it moves, not spatial gradients across its body. Finally, to reach a reasonable signal-to-noise ratio requires that the cell average over time for more than one second.

Why don't the bacteria integrate for longer, and reduce the noise further? If you look closely at the trajectories of the bacteria, you can see that the longer runs curve a bit. In fact, the bacteria are sufficiently small that their own rotational Brownian motion disorients them on a time scale of ten or fifteen seconds. So, if you integrate for longer than this, you are no longer integrating something related to the gradient in a particular direction, or even your current direction of motion. This suggests that there is a physical limit setting the longest useful integration time.

Berg and Purcell also argued that there is a minimum

⁴² I think there is an opportunity for a better experiment here. One could imagine analyzing the moments of transition from run to tumble (and back) in the same way that we analyze the action potentials from sensory neurons (see Section II.C), measuring the reliability of discrimination between small differences in concentration or reconstructing the concentration vs. time along the trajectory of a freely swimming bacterium.

useful integration time. Recall that molecules moving via diffusion traverse a distance $x_{\text{diff}} \sim \sqrt{Dt}$ in a time t ; in contrast, swimming at velocity v moves the bacterium by a distance $x_{\text{swim}} \sim vt$. For short times, diffusion, with its square root dependence on time, goes farther than ballistic swimming motion. This means that on short time scales, the molecules that the bacterium sees along its path are the same molecules, and hence it really isn't combining statistically independent measurements. So, there is a minimum useful integration time (assuming you want to improve the signal-to-noise ratio by integrating) of $\tau \sim D/v^2$, and this works out to be about one second. **Put in a pointer to a problem in the next section.**

So, the strategy of *E. coli* for measuring gradients is incredibly constrained by physics. To reach the observed performance, it has to count nearly every molecule that arrives at its surface. Even with this near ideal behavior, it can work only by making comparisons across time, not space, and estimates of time derivatives have to be averaged for a few seconds, not more and not less. This set of predictions about chemotactic strategy is almost parameter free, even if not precisely quantitative.

What do real bacteria do? We have already seen that they make temporal comparisons. Does the detailed form of these comparisons agree with the Berg-Purcell predictions? Although one could probably do better with modern experimental techniques, the best test was done in the early 1980s. In these experiments, bacteria were tethered to a glass slide and exposed to changing concentrations of attractants or repellents; a long series of such

observations is then combined to measure the probability that the flagellar motor is rotating counterclockwise (corresponding to running) as function of time relative to the changing concentration. A summary of these experiments is shown in Fig 47. We see that the probability of running is modulated by the time derivative of the concentration, averaged over a window of a few seconds, exactly as predicted by the Berg-Purcell argument.

Being sensitive to a derivative means that the response to a step comes back almost exactly to the baseline before the step, as seen at right in Fig 47, so that the constant signal is ignored at times long after it was turned on. This gradual 'forgetting' of a constant signal is common in biological systems, and such phenomena are called 'adaptation.' All of our sensory systems exhibit adaptation, the most familiar being the experience of stepping into a dark movie theater or out into the bright sunlight; at first we are acutely aware of the large difference in overall light intensity, but after a while everything looks normal and we are insensitive to the absolute photon flux. The case of bacteria is interesting because it seems that the adaptation is nearly perfect.

Experiments of the sort pictured in Fig 47 also make it possible to estimate the absolute sensitivity of the system in perhaps more compelling units. **[should put the numbers here, maybe reproduce a figure]** We now know how many receptors there are on the cell's surface, and so we can convert changes in concentration into changes in the number of occupied receptors. Indeed, one extra occupied receptor leads to a significant change in the probability of running vs tumbling. So, as expected, the bacterium is responding to individual molecular events.

This all seems a great success: much of bacterial behavior is understandable, semi-quantitatively, as a response to the physical constraints posed by life at low Reynolds' number and the noise in molecular counting; one can go further and say that bacterial behavior is near optimal in relation to this noise. On the other hand, many questions are left hanging.

First, can we turn the ideas about maximum and minimum useful integration times into a theory of optimal filtering that would predict, quantitatively, the form of the impulse responses in Fig 47? We should be able to do this, but I don't think anyone has really managed to get it right. There have been some serious attempts, but I think the issue still is open. One might also wonder whether it even makes sense to formulate this problem for individual bacteria, as opposed to looking at competition or cooperation in a population; this is related to the question of what, precisely, one thinks is being optimized by the behavior. It seems likely that any theory of optimal strategies will predict that this optimum is context dependent; here we should note that quantitative characterization of chemotactic behavior has not been pursued under a very wide range of stimulus conditions, so we may be missing the data we need to test




FIG. 47 Impulse responses in bacterial chemotaxis, from Block et al (1982). At left, changes in the probability of counterclockwise rotation of the motor, corresponding to running, as a function of time in response to a pulse of attractant (top) or repellent (bottom). We see that the form of the response is equivalent to integrating the time derivative of the input over a window or several seconds. At right, the response to a step of attractant again has the form expected if we integrate the derivative over a short window. The real data are compared with a prediction based on integrating the response to impulses shown at left, and the agreement is good, as if the system were linear.




FIG. 48 Biochemical amplification in the chemotactic response. [Redraw this to make it more obvious that there is a cascade, as in rod photoreceptors.] At left, binding of chemoattractants to their receptors shifts the equilibrium between active and inactive forms of the kinase CheA. At right, the active kinase phosphorylates CheY, and this is balanced by the action of the phosphatase CheZ. CheY~P binds to the flagellar motor and promotes clockwise rotation, which drives tumbling. The motor is extremely sensitive to small changes in the CheY~P concentration; data redrawn from Cluzel et al (2001).

such theories when they emerge.

The second question is about the mechanisms that make possible the extreme sensitivity of chemotaxis. Much progress has been made, although again some issues are open. As with the rod cell, there is a cascade of biochemical events that leads from input (here, binding to receptors on the cell surface) to output (direction of motor rotation). Since input and output are spatially separated, it is not surprising to find that there is an internal signaling molecule that diffuses through the cell. In rods, this is a small molecule (cGMP), but for bacterial chemotaxis it is a protein called CheY. More precisely, this protein can be phosphorylated, and in its phosphorylated form CheY~P it binds to the motor and favors clockwise rotation. The receptor molecules on the cell surface are coupled almost directly to the kinase CheA that phosphorylates CheY, as shown schematically in Fig 48. Working backward from the output, we would like to know how the rotational bias of the motor depends on the concentration of CheY~P.

To measure the bias vs CheY~P, one has to do many tricks. It's relatively easy to measure the bias of the motor, either in experiments where the cell is tethered or where it is laying on a slide and one motor stump is sticking up with a bead attached. To know the concentration of a protein in a single cell, we need to make the protein visible, and so this is done by genetic engineering, replacing the normal CheY with a fusion between this protein

and the green fluorescent protein [put clear discussion of GFP in the first place where it comes up—perhaps here?], and arranging for the expression of this fusion protein to be controlled by signals that can be applied externally. Finally, we need to know the concentration of the phosphorylated form of the protein, and this is very difficult. But once phosphate groups are attached to a protein, they stay there until removed by another enzyme (the phosphatase). So, if we genetically engineer the bacterium to remove the phosphatase, we will surely screw up the overall chemotactic response, but we can then be sure that all the CheY will be in its phosphorylated state. The result of all this is shown in Fig 48.

Problem 57: Absolute concentration measurements. In this problem you should try to understand how Cluzel et al were able to put the CheY~P concentration on an absolute scale. Bacteria can be engineered to make a fluorescent version of many naturally occurring proteins. While the fluorescence signal that we then see under a microscope is proportional to the number of molecules under illumination, it can be difficult to measure the proportionality constant in an independent experiment. One can circumvent this problem by watching small numbers of molecules diffusing randomly in and out of an illuminated volume inside an individual cell and using the *variance* in the fluorescence intensity, along with its mean value, to make an absolute measurement of the concentration of the molecules.⁴³

(a.) Explain (qualitatively) how this measurement might work. What do you gain by using both the variance and the mean of this signal? How can the fluctuating fluorescence signal be analyzed further to give an estimate of the protein diffusion constant?

(b.) Now let's convert the above intuition into a quantitative framework for analysis of the data. Consider the concentration $c(\vec{x}, t)$ of fluorescent molecules at different points in space and time. It fluctuates and the deviation δc of the concentration from its average value \bar{c} is uncorrelated between different points in space (but the same instant of time). Show that the analytic statement

$$\langle \delta c(\vec{x}, t) \delta c(\vec{x}', t) \rangle = \bar{c} \delta(\vec{x} - \vec{x}') \quad (322)$$

of this fact is equivalent to the 'intuitive' remark that the variance of the number of molecules in a volume is equal to the mean number.

(c.) If the system starts with some fluctuation in the concentration $c(\vec{x}, 0) = \bar{c} + \delta c(\vec{x}, 0)$, this profile will relax according to the diffusion equation. Since the diffusion equation is linear, this means that the profile of fluctuations at time t , $\delta c(\vec{x}, t)$, can be written as a linear operator acting on the initial condition $\delta c(\vec{x}, 0)$. Show that this linear relationship can be written as

$$\delta c(\vec{x}, t) = \int d^3y \left(\frac{1}{\sqrt{4\pi Dt}} \right)^3 \exp(-|\vec{x} - \vec{y}|^2/4Dt) \delta c(\vec{y}, 0) \quad (323)$$

where D is the diffusion constant.

⁴³ Some of the ideas in this problem will, admittedly, be clearer after the discussion in the next section. Still, this should be workable now, and may provide a useful introduction to what comes next. This problem was originally designed as part of a general examination for Physics PhD students, written together with Curt Callan.




FIG. 49 A model for the modulation of rotor bias by binding of CheY~P. CheY~P molecules bind independently to multiple sites around a ring. When all sites are empty the equilibrium favors the counterclockwise rotating state. Binding is stronger to the clockwise state, however, so that as more sites are occupied the equilibrium shifts.

(d.) When we bring light to a focus under the microscope, we effectively weight the points around the focus with a Gaussian function, so that the light intensity collected from the fluorescent molecules will be proportional to

$$s(t) = \int d^3x c(\vec{x}, t) \exp(-|\vec{x}|^2/\ell^2) \quad (324)$$

where ℓ is the size of the focal region (roughly the size of the wavelength of light). Using the results above, show that the temporal correlation function of this signal is given by

$$\langle \delta s(t) \delta s(0) \rangle \propto (|t| + \tau)^{-3/2}, \quad (325)$$

and relate the correlation time τ to the diffusion constant D and the size of the focal region ℓ . As a hint, note that in doing the multidimensional Gaussian convolution integrals that show up in the last step of this computation, it is a good idea to do them Cartesian coordinate by Cartesian coordinate. This gives a precise method for extracting the diffusion constant from the fluctuating fluorescence signal.

What we see most clearly from Fig 48 is that the motor is remarkably sensitive to small changes in concentration of CheY~P. One can fit a function of the form

$$P_{cw} = \frac{c^n}{c^n + K^n}, \quad (326)$$

with $K \sim 3 \mu\text{M}$ and $n \sim 10$, although the data are almost within errors of being a step function. “Hill functions” of this form often are interpreted to mean that n molecules bind together and trigger the output that we are measuring; these and other ideas about the cooperative response of biological molecules are reviewed in Appendix A.4.

In this case it might make more sense to think about a model as in Fig 49, which is a version of the Monod–Wyman–Changeux model for cooperativity. Here we

imagine multiple binding sites arrayed around a ring. CheY~P molecules bind independently to each site, but the strength of the binding depends on whether the whole structure is rotating clockwise or counterclockwise. Qualitatively, if binding is stronger in the clockwise state, then increasing the concentration of CheY~P will shift the equilibrium toward the clockwise state.

Quantitatively, we can work out the predictions of the model in Fig 49 using statistical mechanics, on the hypothesis that all the binding events and the structural transitions of the motor between clockwise and counterclockwise states come to equilibrium. One might worry about the latter assumption—after all, if the motor were truly at equilibrium it wouldn’t be rotating and generating force—but let’s proceed. Consider one possible state of the system, say clockwise rotation with m out of the n sites filled by CheY~P molecules. We need to assign this state a weight in the Boltzmann distribution. We can assume that the clockwise state has an intrinsic (free) energy E_{cw} . With k molecules bound, the energy is lowered by mF_{cw} , where F_{cw} is the binding energy in the clockwise state, but we also had to take these k molecules out of solution, and this shifts the free energy by m times the chemical potential, $m\mu = mk_B T \ln(c/c_0)$, where c is the concentration of CheY~P and c_0 is a reference concentration. Finally, since the m occupied sites could chosen out of the n possibilities in many ways, there is a combinatorial factor. Putting these terms together we have

$$\binom{n}{m} \exp \left[-\frac{1}{k_B T} (E_{cw} - mF_{cw} - mk_B T \ln(c/c_0)) \right] = \binom{n}{m} \left(\frac{c}{K_{cw}} \right)^m e^{-E_{cw}/k_B T},$$

where $K_{cw} = c_0 e^{-F_{cw}/k_B T}$. To compute the probability of being in the clockwise state we have to sum over all the different occupancies, and normalize by the partition function, which includes a sum over the counterclockwise states:

$$P_{cw} = \frac{1}{Z} \sum_{m=0}^n \binom{n}{m} \left(\frac{c}{K_{cw}} \right)^m e^{-E_{cw}/k_B T} \quad (327)$$

$$= \frac{1}{Z} e^{-E_{cw}/k_B T} (1 + c/K_{cw})^n, \quad (328)$$

where

$$Z = e^{-E_{cw}/k_B T} (1 + c/K_{cw})^n + e^{-E_{ccw}/k_B T} (1 + c/K_{ccw})^n. \quad (329)$$

We can put this result in a more compact form,

$$P_{cw} = \frac{1}{1 + \exp[\theta - g(c)]} \quad (330)$$

$$\theta = (E_{cw} - E_{ccw})/k_B T \quad (331)$$

$$g(c) = n \ln \left(\frac{1 + c/K_{cw}}{1 + c/K_{ccw}} \right). \quad (332)$$

Notice that if $K_{cw} \ll c \ll K_{ccw}$, then this becomes the Hill function in Eq (326).

Problem 58: MWC model of rotor bias. Explore the parameter space of the model we have just described. Are there regimes, other than $K_{cw} \ll c \ll K_{ccw}$, where one can reproduce the steep dependence of P_{cw} on c observed by Cluzel et al (2001)? Keep in mind that the actual number of binding sites n could be very large.

So part of the answer to how the the bacterium is so sensitive to small changes in the external concentration of attractants or repellents is that the motor is very sensitive to small changes in the concentration of CheY~P. This is not implausible, since the structure of the motor (which is complicated) suggests locations for as many as $n = 34$ sites where CheY~P could bind around a ring of radius $R \sim 45$ nm.

Having such strong sensitivity to the CheY~P concentration means that, in roughly the one second it takes for the motor to switch once, one can be sure whether the concentration was $\delta c/c \sim 1/n \sim 10\%$ above or below the critical value $c = K$. But from Berg and Purcell we might expect that there is a limit on this precision set by random arrival of the CheY~P molecules at the motor, and this should be $\delta c/c \sim 1/\sqrt{DRc\tau_{avg}}$, treating the whole motor ring as one big receptor. With diffusion constants for proteins, including CheY, in the range of $D \sim 1\mu\text{m}^2/\text{s}$, this suggests that the limit with one second of integration is not much smaller than 10% (see more details in the next lecture). So, cooperative action of many signaling molecules generates a steep slope, but the system still has to suppress other sources of noise since even this last step in the cascade of events is operating close to the fundamental limits set by noise considerations.

The observations on the sensitivity of the motor tell us that the bacterium can generate a significant response even from a small fractional change in the concentration of CheY~P. Still, we need to understand the biochemical processes that lead from essentially single molecular events to these quasi-macroscopic changes in molecule number.⁴⁴ [Probably want to say a few words about the sources of gain: activity of CheA*, and the cooperativity among receptors that allows one ligand to activate many CheAs. Need to learn more about the numbers here. Might be nice to compare MWC-style model of motor with MWC-style model of receptors. At the end of the day, is this similar to the rod cell or not? Can we conclude that we understand the gain?]

⁴⁴ At $c \sim 3\mu\text{M}$, a cell with volume $\sim 1\mu\text{m}^3$ has ~ 2000 molecules of CheY~P, so even a ten percent change in concentration involves hundreds of molecules.

Even if we consider the origins of gain to be understood, there is a major problem. Figure 48 shows that extreme sensitivity must coexist with a very tight regulation, since if the concentration of CheY~P drifts far away from $c \sim K$, the cell loses all sensitivity to changes. This combination of sensitivity to small changes without accumulation of large variations poses significant problems, which we will take up in the next Chapter.

The last of the major questions left open by the Berg–Purcell analysis is whether we do a full, honest calculation that leads to the their limit on the precision of concentration sensing? What Berg and Purcell wrote down makes absolutely no reference to the messy details of what actually happens to molecules as they are counted. This could be wonderful, because it would mean that can say something about the limits to precision in *all* biochemical signaling systems, regardless of details. Alternatively, the absence of details might be a disaster, a clue that we have simply missed the point.

As mentioned at the start of this section, chemical signaling in ubiquitous in biological systems, and chemotaxis provided us with one clear example where we could think about the limits to counting molecules. We would like to know if these limits can be made rigorous, and




FIG. 50 Control of gene expression by transcription factors. Synthesis of a protein involves transcription of the DNA coding for that particular protein, and translation of the resulting mRNA. An important component of control in these systems is the binding of transcription factors to the DNA, at specific sites near the start of transcription, in the promoter or enhancer region. Transcription factors are themselves proteins, so this regulatory process naturally leads to a network of interactions; here we focus, for simplicity, on one input (the concentration of the transcription factor) and one output (the concentration of protein #1). Note that in bacteria all of this happens in one compartment, while in eukaryotic cells the DNA is in the nucleus and mRNA is transported out to the cytoplasm, where translation occurs. Nothing in this figure is to scale. [redraw figure to get rid of the network, which here is a distraction]

if they can be applied to processes that occur inside of cells, rather than just to the sensing of external signals as in chemotaxis. To see what is at stake, let's think about the regulation of gene expression (Fig 50). We recall that every cell in our bodies has the same DNA. What makes a liver cell different from a neuron in your brain is that it reads out or “expresses” different genes, making different proteins. Importantly, this is not just a discrete choice made once in your lifetime. Given that certain proteins are being made, the numbers of these molecules are constantly adjusted to match the needs of the cell. This happens also in bacteria, which adjust, for example, the concentrations of the enzymes needed to metabolize different nutrients that might or might not be present in the environment; much of what we know about the regulation of gene expression has its roots in work on this sort of metabolic control in bacteria.

There are many ways in which gene expression is controlled. As a simple example, note that if we want to regulate the number of proteins in the cell we can change either the rate at which they are made or the rate at which they are degraded, and both of these things happen. The synthesis of a protein involves two very different steps, transcription from DNA to messenger RNA and translation from mRNA to protein, and again there is regulation of both processes. All this being said, we will focus our attention on the regulation of transcription, that is the reading of the DNA template to make mRNA.⁴⁵

In order to make mRNA, a complex of proteins (including the RNA polymerase) must bind to the DNA and ‘walk’ along it, spewing out the mRNA polymer as it walks. In order for all of this to happen, the RNA polymerase has to find the right starting point. One can imagine that this can be inhibited simply by having other proteins bind to nearby sites along the DNA. Alternatively, binding of proteins to slightly different positions near the starting point could help the RNA polymerase to find its way. Both of these things happen: proteins called transcription factors can act both as repressors and as activators of mRNA synthesis. The key step in this regulation is thought to be the binding of the transcription factors to specific sites near the RNA polymerase start site, as schematized in Fig 50; the whole segment of DNA involved in the control and initiation of transcription is called the “promoter.” In higher organisms, the regions involved in regulation can be very large indeed, and usually are called “enhancers” to avoid conjuring the simplified image in Fig 50, which is more literally applicable in bacteria. Binding sites are specific because the transcription factor protein is selective for particular DNA sequences, and much can be said about the na-

ture of this specificity. For now the important point is that such regulatory systems are, in effect, sensors of the transcription factor concentration.

Problem 59: Autoregulation. Perhaps the simplest model of transcriptional regulation is one in which a gene regulates its own expression. Let the concentration (or number of molecules) of the protein be g , and assume that n of these molecules bind cooperatively to the promoter region of the gene. If the binding activates expression, and proteins are degraded in a simple first order process with lifetime τ , then it is plausible that the dynamics of g are given by

$$\frac{dg}{dt} = r_{\max} \frac{g^n}{g^n + g_{1/2}^n} - \frac{g}{\tau}. \quad (333)$$

(a.) Explain the significance of the parameters r_{\max} , n and $g_{1/2}$. Show that there is a range of these parameters in which the system is bistable. More precisely, show that you can find three steady states, and that two of these are stable and one is unstable. What are the time constants for relaxation to these steady states? How do these times compare with the lifetime τ of the protein?

(b.) Really the protein binding regulates the synthesis of mRNA, which in turn is translated by the ribosomes into protein. If m is the mRNA concentration (or number of molecules), then a plausible set of equations is

$$\frac{dm}{dt} = e_{\max} \frac{g^n}{g^n + g_{1/2}^n} - \frac{m}{\tau_m} \quad (334)$$

$$\frac{dg}{dt} = r_{\text{trans}} m - \frac{g}{\tau_p}, \quad (335)$$

where e_{\max} is the maximal transcription (“expression”) rate, r_{trans} is the rate at which mRNA molecules are translated into protein, and the lifetimes of protein and mRNA are τ_p and τ_m , respectively. Under what conditions will this more complete model be well approximated by the simpler model above? Are the steady states of the two models actually different? What about their stability?

(c.) Suppose that instead of activating its own expression, the protein acts as a repressor of its own expression. Find the analog of Eq (333) in this case and show that there is only one steady state, and that this state is stable.

(d.) Expand your discussion of the auto-repressor to include the mRNA concentration, as in Eq's (334, 335). Find the steady state and linearize the equations around this point. Do you find exponential relaxation toward the steady state for all values of the parameters? Is it possible for the steady state to become unstable? Explain qualitatively what is happening, and go as far as you can in analyzing the situation analytically.

The binding sites along DNA for the transcription factors have linear dimensions measured in nanometers, perhaps $a \sim 3$ nm. The diffusion constants of proteins in the interior of cells is in the range of $D \sim 1 \mu\text{m}^2/\text{s}$. Many transcription factors act at nanoMolar concentrations, and it is useful to note that $1 \text{ nM} = 0.6 \text{ molecules}/\mu\text{m}^3$. Putting these together we have $Dac \sim 1.8 \times 10^{-3} \text{ s}^{-1}$. Thus, the Berg–Purcell limit predicts that the smallest

⁴⁵ For a bit about the basics of DNA structure, see Appendix A.5.

changes in transcription factor that can be reliably detected are

$$\frac{\delta c}{c} \sim \frac{1}{\sqrt{Dac\tau_{\text{avg}}}} \sim \sqrt{\frac{10 \text{ min}}{\tau_{\text{avg}}}}. \quad (336)$$

Taken at face value, this suggests that truly quantitative responses—say, to 10% changes in transcription factor concentration—would require hours of integration. This is seldom plausible.

One should not take this rough estimate too literally. I think the message is not the exact value of the limiting precision, but rather that once concentrations fall to the nM range, small changes will be very hard to detect. If cells do detect these small changes, then almost certainly they will be bumping up against the physical limits set by counting molecules, assuming that Berg and Purcell give us a good estimate of these limits. So, this is what we need to check.

In Appendix A.6, we look in detail at how to make the Berg–Purcell limit more rigorous. The key idea is that fluctuations in concentration, and in many examples of binding to receptor sites, represent fluctuations in thermal equilibrium, and thus are susceptible to the same analyses as Brownian motion, Johnson noise, and other examples of thermal noise. These analyses show how one can separate the limiting noise level from the extra noise that is associated with all the biochemical complexities which Berg and Purcell ignored. The result, then is that the Berg–Purcell argument can be made rigorous, both for single receptors and for arrays of receptors, and their simple formula gives us a *lower bound* on the noise in biochemical signaling. This is important because, as noted at the start of this discussion, the Berg–Purcell limit doesn’t make reference to any of the detailed biochemistry of what happens when the signaling molecules bind to their targets. Rather, the limit depends on the physical nature of the signal itself. The fact that we can make the Berg–Purcell argument rigorous encourages us to look more broadly and see if there are other cases in which biological systems approach these physical limits to their signaling performance.

Would like to discuss chemotaxis in larger cells—neutrophils, Dictyostelium, ...

Another important example of chemotaxis occurs during the development of the brain. Individual neurons start as relatively compact cells, and then extend their axons to find the other cells with which they must make synapses. This processes is guided by gradients in a variety of signaling molecules. Although there are many beautiful observations on these phenomena in vivo, it is not so easy to do a controlled experiment where one allows cells to migrate in well defined gradients. One approach to this is shown in Fig [reproduce figures from Rosoff et al], where cells grow in a collagen matrix that is “printed” with droplets of growth factor at varying

densities. Relatively quickly, diffusion acts to smear the rows of drops into a continuous gradient, which can be directly observed when the molecules are labelled with fluorophores. These measurements also allow an inference of the diffusion constant in this medium, $D \sim 8 \times 10^{-7} \text{ cm}^2/\text{s}$. The growth cones which guide the axon have linear dimensions $a \sim 10 \mu\text{m}$, and these experiments found that sensitivity to gradients is actually maximal in a concentration range near $c \sim 1 \text{ nM}$. Under these conditions, then, we have $Dac \sim 500 \text{ s}^{-1}$. Quite astonishingly, however, the cells seem to grow differentially in the direction of gradients that correspond to concentration differences of order one part in one thousand across the diameter of the growth cone. In order for this signal to be above the Berg–Purcell limit on the noise level, the cell must integrate for $\tau_{\text{avg}} \sim 2000 \text{ s}$, a reasonable fraction of an hour.

In truth, we don’t know the time scale over which growth cones are integrating as they decide which way to turn, even in the more controlled in vitro experiments. We do know that the pace of neural development is slow—hours to days rather than minutes. Qualitative aspects of axonal behavior are consistent with the idea that the time scales of their movements are determined by the need to integrate long enough to generate reliable directional signals, from the rapid “exploration” by cellular appendages to the dramatic slowing down near critical decision points, such as the optic chiasm where the axons of ganglion cells emerging from the retina must decide whether to go toward the right or left half of the brain.⁴⁶ It is attractive to think that the reliability with which cells in our brain find their targets is set by such basic physical principles, but we don’t quite have enough data to say this with certainty.

Let us return to the problem that motivated our search for generality, the transcriptional regulation of gene expression. Until the last decade, there were essentially no direct measurements on the reliability of such regulatory mechanisms. Before we look at the new data, though, we need one more set of theoretical ideas.

Proteins are synthesized and degraded, and the simplest assumption is that these are single kinetic steps. Suppose we start just with synthesis, at some rate s molecules per second. We have seen that rate constants should be interpreted as the probability per unit time for individual molecular events. Thus, if we ask about the probability of finding exactly N molecules in the system at time t , this probability $P(N; t)$ obeys the “master

⁴⁶ At these decision points it seems likely that the cells must reach rather high signal-to-noise ratios, since the error probabilities are small. [can we say something quantitative here?]

equation”

$$\frac{\partial P(N; t)}{\partial t} = sP(N-1; t) - sP(N; t), \quad (337)$$

except of course at $N = 0$ where we have

$$\frac{\partial P(0; t)}{\partial t} = -sP(0; t). \quad (338)$$

We can solve these equations iteratively. We start with no molecules, so $P(0, 0) = 1$, while $P(N \neq 0, 0) = 0$. Then Eq (338) tells us that

$$P(0, t) = e^{-st}. \quad (339)$$

If we substitute into Eq (337) for $P(1, t)$, we have

$$\frac{\partial P(1; t)}{\partial t} = -sP(1; t) + sP(0; t) \quad (340)$$

$$\Rightarrow P(1, t) = \int_0^t dt' e^{-s(t-t')} sP(0; t') \quad (341)$$

$$= \int_0^t dt' e^{-s(t-t')} s e^{-st} \quad (342)$$

$$= s e^{-st} \int_0^t dt' = e^{-st} (st). \quad (343)$$

We can go through the same calculation for $P(2; t)$:

$$P(2; t) = \int_0^t dt' e^{-s(t-t')} sP(1; t') \quad (344)$$

$$= e^{-st} \int_0^t dt' s^2 t' \quad (345)$$

$$= e^{-st} \frac{(st)^2}{2}. \quad (346)$$

This suggests that, for all N ,

$$P(N; t) = e^{-st} \frac{(st)^N}{N!} \quad (347)$$

where k is the probability per unit time for the decay of one molecule.

Now it is possible for the synthesis and degradation reactions to balance, generating a steady state. In this steady state the distribution of the number of molecules must obey

$$0 = sP(N-1) - (s+kN)P(N) + k(N+1)P(N+1). \quad (350)$$

Problem 60: Checking the Poisson solution. Verify that Eq (347) solves the master equation describing a single synthesis reaction at rate s , Eq (337).

Equation (347) is telling us that, as the synthesis reaction proceeds, the number of molecules that has been synthesized obeys the Poisson distribution. From what we have said about the Poisson distribution in the discussion of photon counting (Section I.A and Appendix A.1), you should recognize that the mean number of molecules is

$$\langle N \rangle \equiv \sum_{N=0}^{\infty} NP(N; t) = st, \quad (348)$$

which makes perfect sense. Further, the variance in the number of molecules is equal to the mean, at all times.

[This discussion is written without any figures. Maybe we need some schematics?]

What happens when we add degradation to this picture? Now the state of the system can change in several ways, all of which will modify the probability that there are exactly N molecules. First, synthesis can cause the N molecules to become $N+1$, reducing $P(N, t)$. Second, we can have the transition from $N-1$ to N molecules, which increases $P(N, t)$. Note that these first two terms were already present in our simpler model. The third process is where degradation takes N molecules and eliminates one, resulting in $N-1$ molecules. Since each molecule makes its transitions independently, the rate of this process must be proportional to N , and this reduces $P(N, t)$. Finally, if there were $N+1$ molecules, degradation results in N , increasing $P(N, t)$; again because each molecule is independent, the rate of this process must be proportional to $N+1$. Putting the terms together we have

$$\frac{\partial P(N; t)}{\partial t} = -sP(N; t) + sP(N-1; t) - kNP(N; t) + k(N+1)P(N+1; t), \quad (349)$$

To solve this equation it is useful to regroup the terms,

$$-sP(N-1) + kNP(N) = -sP(N) + k(N+1)P(N+1). \quad (351)$$

where the left hand side now refers to the forward and backward rates between states with $N-1$ and N molecules, while the right hand side refers to the transitions between N and $N+1$. All that we require is that the

two sides be equal, but suppose we try to set each side separately to zero, which corresponds to “detailed balance” among the transitions into and out of each state. Then from the left hand side we have

$$\frac{P(N)}{P(N-1)} = \frac{s}{kN}, \quad (352)$$

while from the right we have

$$\frac{P(N+1)}{P(N)} = \frac{s}{k(N+1)}. \quad (353)$$

But except for $N \rightarrow N+1$, these are the same equation. Thus, the steady state of this system does obey detailed balance, and we can solve by iterating Eq (352):

$$P(1) = \frac{s}{k}P(0) \quad (354)$$

$$P(2) = \frac{s}{2k}P(1) = \frac{(s/k)^2}{2}P(0) \quad (355)$$

$$P(3) = \frac{s}{3k}P(2) = \frac{(s/k)^3}{3!}P(0), \quad (356)$$

and, in general,

$$P(N) = \frac{(s/k)^N}{N!}P(0). \quad (357)$$

Finally we can fix the value of $P(0)$ by insisting that the distribution be normalized, and we find

$$P(N) = e^{-M} \frac{M^N}{N!}, \quad (358)$$

which again is the Poisson distribution, with mean $M = s/k$.

Problem 61: The diffusion approximation. If N is not too small we expect that $P(N; t)$ and $P(N \pm 1; t)$ are not too different. Thus we should be able to approximate using a Taylor series,

$$P(N \pm 1; t) \approx P(N; t) \pm \frac{\partial P(N; t)}{\partial N} + \frac{1}{2} \frac{\partial^2 P(N; t)}{\partial N^2}. \quad (359)$$

(a.) Show that this approximation turns the master equation in Eq (349) into something that looks more like the diffusion equation. What is the effective potential in which the “coordinate” N is diffusing?

(b.) Why does it make sense to stop the Taylor series after two derivatives? What happens if we stop after one?

(c.) How does the steady state solution that you obtain in the diffusion approximation compare with the exact solution (the Poisson distribution)?

Problem 62: Langevin equations for chemical kinetics. We know, as reviewed in Section II.A, that we can describe Brownian motion by either a diffusion equation or a Langevin equation. In more detail, we started with kinetics that, in the macroscopic limit, correspond to the dynamics

$$\frac{dN(t)}{dt} = s - kN(t). \quad (360)$$

We would like to describe the noisy version of these dynamics as

$$\frac{dN(t)}{dt} = s - kN(t) + \zeta(t), \quad (361)$$

where—inspired by the Brownian motion example—we expect that the noise $\zeta(t)$ is white, but the strength might depend on the state of the system, so that

$$\langle \zeta(t)\zeta(t') \rangle = T_{\text{eff}}[N(t)]\delta(t-t'), \quad (362)$$

where to remind us of the analogy to Brownian motion we can refer to the noise strength as an effective temperature T_{eff} .

(a.) Find the effective temperature that will reproduce the diffusion equation that you derived in the preceding problem.

(b.) If we integrate Eq (361) over a very small time interval $\Delta\tau$, we obtain

$$\Delta N \equiv N(t + \Delta\tau) - N(t) \quad (363)$$

$$= [s - kN(t)]\Delta\tau + \int_0^{\Delta\tau} dt' \zeta(t + t'). \quad (364)$$

But if $\Delta\tau$ is small enough, we know that the changes in the number of molecules should be $\Delta N = 0$ or $\Delta N = \pm 1$. Going back to the master equation [Eq 349], identify these transition probabilities. From these probabilities, show that the mean change in the number of molecules is the first term in Eq (364), $\langle \Delta N \rangle = [s - kN(t)]\Delta\tau$. Continuing, show that the variance in ΔN is given by $\langle (\delta\Delta N)^2 \rangle = [s + kN(t)]\Delta\tau$.

(c.) To reproduce the variance in ΔN , we must have

$$\left\langle \left(\int_0^{\Delta\tau} dt' \zeta(t + t') \right)^2 \right\rangle = [s + kN(t)]\Delta\tau. \quad (365)$$

Use this, together with Eq (362), to show that

$$T_{\text{eff}}[N(t)] = s + kN(t). \quad (366)$$

Does this agree with your result in (a.)?

So, these simplest of kinetic schemes for the synthesis and degradation of molecules predict that the distribution of the number of molecules (“copy numbers”) should be Poisson. Certainly we can imagine kinetic schemes for which the fluctuations in copy number will be larger than Poisson. For example, if the simple picture of synthesis and degradation were correct for messenger RNA, but each mRNA leads to the synthesis of b proteins, then the mean number of proteins will be larger than the mean number of mRNA molecules by this factor b , $\langle N_p \rangle = b\langle N_{\text{mRNA}} \rangle$, but the variance will be larger by a factor of b^2 , $\langle (\delta N_p)^2 \rangle = b\langle (\delta N_{\text{mRNA}})^2 \rangle$. Thus, if we count protein molecules, the variance will be larger than the mean, $\langle (\delta N_p)^2 \rangle = b\langle N_p \rangle$, and hence the protein copy numbers are more variable than expected from the Poisson distribution. Notice that this is true even though we have assumed that the translation from mRNA to protein is completely noiseless, with each mRNA making exactly b proteins. Variance beyond the Poisson expectation here arises simply from amplification. This is exactly the same argument made about photons and spikes from ganglion cells in the retina, in Section I.D.

With this background, what can we measure? Counting protein molecules is not easy. Over the last decades,




FIG. 51 Noise in the regulation of gene expression, from Elowitz et al (2002). A population of *E coli* express two fluorescent proteins of different colors, CFP and YFP, both under the control of the *lac* repressor. At left, expression is repressed, copy numbers are low, and color variations are substantial. Thus, although the two genes see the same regulatory signals, there is intrinsic variation in the output. At right, repression is relieved, expression levels are higher, and color variations are substantially smaller.

we have seen a huge improvement in the methods of optical microscopy, to the point where we can literally see the light emitted from a single fluorescent molecule. But most biological molecules, and most proteins in particular, are not fluorescent. Indeed, until relatively recently the only proteins with interesting spectroscopic signatures in the visible part of the spectrum (e.g., the visual pigments and the heme proteins) involved a smaller molecular cofactor bound to the protein (retinal, heme). These cofactors are synthesized by separate, often complex pathways. Thus while it might be possible to engineer a cell to make a pigment protein just by splicing the relevant gene into its genome, it would be almost impossible to introduce the entire synthetic machinery for the cofactor. This is why the discovery of the green fluorescent protein in a species of jellyfish turned out to be so important. In contrast to the proteins which require cofactors for their fluorescence, these molecules are intrinsically fluorescent [Need a figure showing structure, point to why this is possible, etc.. Maybe this discussion should come earlier?] Since the isolation of the original GFP, many variants have been synthesized, in a variety of colors.

The simplest experiment to probe noise in the expression of a gene would be to introduce the gene for GFP into a bacterium, and just look at the levels of fluorescence—the brightness will be proportional to the number of molecules, and with luck we can even calibrate the proportionality factor. But expression levels could vary for uninteresting reasons. Cells vary in size

as they grow and divide. There can be variations in the number of ribosomes, which will change the efficiency of translation but it probably doesn't make sense to call these variations “noise.” How do we separate all these different sources of variation from genuine stochasticity in the processes of transcription and translation?

If we go back to Fig 50, we see that the transcription of a gene into RNA is controlled by the binding of transcription factor proteins to a segment of DNA called the promoter or (in higher organisms) enhancer region. Suppose that we make two copies of the same promoter, put one next to the gene for a green fluorescent protein and one next to the gene for a red fluorescent protein, and then reinsert both of these into the genome. Now all variations in the state of the cell that affect the overall efficiency of transcription and translation will change the levels of green and red proteins equally. If the regulatory signals were noiseless, and the independent processes of transcription and translation of the two proteins were similarly deterministic, then every cell would be perfectly yellow, having made equal amounts of green and red protein; cells might differ in their total brightness, but the balance of red and green would be perfect. On the other hand, if there really is noise in transcription and translation, or their regulation, then the balance of red and green will be imperfect, and if we look a population of genetically identical cells they will vary in color as well as in brightness.

Figure 51 shows that our qualitative expectations for a “two color” experiment are borne out in real experiments on *E coli*, although “red” and “green” are actually yellow and cyan. In this experiment, the two fluorescent proteins are under the control of the *lac* promoter. In the native bacterium, this promoter controls the expression of enzymes needed for the metabolism of lactose, and if there is a better source of carbon available (or if lactose itself is absent) the bacteria don't want to make these enzymes. There is a transcription factor protein called lac repressor which binds to the *lac* promoter and blocks transcription. By changing environmental conditions, one can tap into the signals that normally tell the bacterium that it is time to turn on the *lac*-related enzymes, and turn off the repression by inactivating the repressor proteins. Thus, not only can we get *E coli* to make two colors of fluorescent protein, we can even arrange things so that we have control over the mean number of proteins that will be made. Everything that we have said thus far about noise in synthesis and degradation reactions predicts that if the cell makes more protein on average, then the fractional variance in how much protein is made should be reduced, and this is exactly what we see in Fig 51.

More quantitatively, in Fig 52 we see the decomposition of the variations into an “extrinsic” part that changes the two colors equally and an “intrinsic” part that corresponds to relative variations in the expression

of the two proteins that are under nominally identical control. If synthesis and degradation of proteins were a Poisson process, then we expect from above that the variance would be equal to the mean; amplification of Poisson fluctuations in mRNA count would leave the variance proportional to the mean. Even if the Poisson model is exact, if we can't calibrate the fluorescence intensity to literally count the molecules, again all we could say is that the variance of what we measure will be proportional to the mean. In fact, the data are described well by

$$\frac{\langle(\delta F)^2\rangle}{\langle F\rangle} = \frac{A}{\langle F\rangle} + B, \quad (367)$$

where the fluorescence is normalized so that the mean under conditions of maximal expression is one, and $A = 7 \times 10^{-4}$ and $B = 3 \times 10^{-3}$. If $B \rightarrow 0$, this is exactly the prediction of the Poisson model, and indeed B is small. Importantly, we can see the decrease in the fractional noise level with the increase in the mean. The absolute numbers also are interesting, since they tell us that cells can—at least under some conditions—set the expression level of a protein to an accuracy of better than 10%.

It has been appreciated for decades that the initial steps in the development of embryos provides an excellent laboratory in which to study the regulation of gene expression. As we have mentioned several times, what makes the different cells in our body different is, funda-




FIG. 52 Separating intrinsic and extrinsic noise, from Elowitz et al (2002). At left, a scatter plot of the fluorescence from the two different proteins show the decomposition into variations in the overall efficiency of transcription and translation (“extrinsic” noise) and fluctuations that change the two expression levels independently (“intrinsic” noise). At right, while the total variance has no simple dependence on the mean expression level, the intrinsic noise goes down systematically as the mean expression level goes up. Quantitatively, we plot the standard deviation σ in fluorescence level, divided by the mean m , as a function of the mean. The dotted line is from Eq (367).

mentally, that they express different proteins. These differences in expression have a multitude of consequences, but the first step in making a cell commit to being one type or another is to turn on (and off) the expression of the correct set of genes. At the start, an embryo is just one cell, and through the first several rounds of cell division it is plausible that the daughter cells remain identical. At some point, however, differences arise, and these are the first steps on the path to differentiation, or specialization of the cells for different tasks in the adult organism.

A much studied example of embryonic development is the fruit fly *Drosophila melanogaster*. We will learn much more about this system in Section III.C, but for now the key point is that in making the egg, the mother sets the initial conditions for development in part by placing the mRNA for key proteins—referred to as the “primary morphogens”—at cardinal points in the embryo. As these messages are translated, the resulting proteins diffuse through the embryo, and act as transcription factors, activating the expression of other genes. An example is Bicoid, for which the mRNA is localized at the (eventual) head; the diffusion and degradation of the Bicoid (Bcd) protein leads to a spatial gradient in its concentration, and we can visualize this by fixing and stain-




FIG. 53 Bicoid (Bcd) and Hunchback (Hb) in the early *Drosophila* embryo. At top, an electron micrograph of the embryo in cell cycle fourteen, with thousands of cells in a single layer at the surface (image courtesy of EF Wieschaus). At the bottom left, the embryo has been exposed to antibodies against the proteins Bcd and Hb, and these antibodies in turn have been labelled by green and red fluorophores, respectively; the fluorescence intensity should be proportional to the protein concentration, perhaps with some background. Bicoid is a transcription factor that activates the expression of Hunchback, and at the bottom right we see a scatter plot of the output (Hb) vs input (Bcd), where each point represents the state of one nucleus from the images at left; from Gregor et al (2007b).

ing the embryo with fluorescent antibodies, as shown in Fig 53. A more modern approach is to fuse the gene for Bcd with a fluorescent protein and substitute this for the original gene; if one can verify that the fusion protein replaces the function of the original, quantitatively, then we can measure the spatial profile of Bcd in a live embryo. Among other things, this approach makes it possible to demonstrate that the fluorescence signal from antibody staining really is proportional to the protein concentration, so we can interpret the data from images such as those in Fig 53 quantitatively.

From our point of view, in constructing the embryo, the mother has created an ideal experimental chamber. After just a few hours, there are thousands of cells in a controlled environment, exposed to a range of input transcription factor concentrations that we can literally read out along the embryo. We can also measure the response to these inputs, for example the expression of the protein Hunchback shown in Fig 53. In fact the targets of Bcd are themselves transcription factors, so conveniently they localize back to the nucleus, and hence each nucleus gives us one data point for characterizing the input/output relation. Taking seriously the linearity of antibody staining we can plot the input/output relation between Bcd and Hb in appropriately normalized coordinates, as in Fig 54, and we can measure the noise in expression by computing the variance across the many nuclei that experience essentially the same input Bcd level.

The first thing we see from Fig 54 is that, consistent with the results from bacteria in Fig 52, the embryo can regulate the expression of Hunchback to $\sim 10\%$ accuracy or better across much of the relevant dynamic range. How does this compare with the physical limits? To measure the reliability of Hunchback's response to Bicoid, we should refer the noise in expression back to the input—if we want to change the output by an amount that is equal to one standard deviation in the noise, how much do we have to change the input? The answer is given by propagating the variance backwards through the input/output relation,

$$\langle(\delta\text{Hb})^2\rangle = \left|\frac{d(\text{Hb})}{d\ln c}\right|^2 \left(\frac{\delta c}{c}\right)_{\text{eff}}^2, \quad (368)$$

where c is the concentration of Bcd, and $(\delta c/c)_{\text{eff}}$ defined in this way should be comparable to the Berg–Purcell limit. In Fig 54 we see that this effective noise level drops down to $(\delta c/c)_{\text{eff}} \sim 0.1$, so the system seems able to respond reliably to $\sim 10\%$ differences in concentration of the input transcription factor.

We have seen, in Eq (336) and the surrounding discussion, that responding reliably to 10% differences in transcription factor concentrations would be very difficult to detect, requiring hours of integration to push the noise level down to manageable levels. This seems generally implausible, but in the fly embryo it is impossible,




FIG. 54 Input/output and noise in the transformation from Bcd to Hb, from Gregor et al (2007b). (A) The input/output relation can be obtained starting from the scatter plot in Fig 53, normalizing the fluorescence intensities as relative concentrations, and then averaging the output Hb expression level across all nuclei that have essentially the same input Bcd level. Blue curves show results for several individual embryos, and red circles with error bars show the mean and standard deviation of Hb expression level vs Bcd input for a single embryo. The inset shows that these data are well fit by a Hill relation [see the discussion around Eq (326)] with $n = 5$ (in red), and substantially less well fit by $n = 3$ or $n = 7$ (in green). (B) The standard deviation of Hb output, measured across the multiple nuclei with the same Bcd input in single embryos; different curves correspond to different individual embryos. (C) Combining the input/output relation and noise levels, we obtain the effective noise level referred to the input, as in Eq (368); blue points are raw data, green line is an estimate of measurement noise, and red circles are the results of subtracting the measurement noise variance, with error bars computed across nine embryos. (D) Correlations in Hb expression noise in different nuclei, as a function of distance.

since the whole process from laying the egg to the establishment of the basic body plan (several steps beyond the expression of Hunchback) is complete within three hours or less. This apparent paradox depends on estimating some key parameters, but in the Bcd/Hb system these can be measured, and the solution to the problem does not seem to lie here.

Problem 63: Effective diffusion constants. Add a problem about the renormalization of diffusion constants by transient binding ... connect to noise levels, in a somewhat open ended second part.

On the other hand, the fly embryo is unusual in that, for much of its early development there are no walls between the cells. Thus, Hunchback mRNA synthesized

in one nucleus will be exported to the neighboring cytoplasm, and the translated protein should be free to diffuse to other nuclei. Thus the Hunchback level in one nucleus should reflect an average over the Bcd signals from many cells in the neighborhood. If Hb has a diffusion constant similar to that of Bcd, then in a few minutes the molecules can cover a region which includes ~ 50 nuclei, and averaging over 50 independent Bcd signals is enough to convert the required integration time from hours to minutes. If this scenario is correct, there should be correlations among the Hb expression noise in nearby nuclei, and this is what we see in Fig 54D. Indeed, the correlation length of the fluctuations is just what we need in order to span the minutes/hours discrepancy. These results suggest strongly that the reliability of the Hunchback response to Bicoid is barely consistent with the physical limits, but only because of spatial averaging.

Can we give a fuller analysis of noise in the Bcd/Hb system? In particular, we see from Fig 54B that the noise level has a very characteristic dependence on the input concentration, which we can also replot vs the mean output, as in Fig 55. This is an interesting way to look at the data, because in the limit where the Poisson noise of synthesis and degradation is dominant we should have

$$\langle(\delta\text{Hb})^2\rangle_{\text{Poisson}} = \alpha\langle\text{Hb}\rangle, \quad (369)$$

where the constant α depends on the units in which we measure expression, but reflects the absolute number of independent molecules that are being made. On the other hand, if the random arrival of transcription factors at their target is dominant, we should have Eq (368) with the effective noise given by the Berg–Purcell limit,

so that

$$\langle(\delta\text{Hb})^2\rangle_{\text{BP}} = \left| \frac{d\langle\text{Hb}\rangle}{d\ln c} \right|^2 \cdot \frac{1}{N_{\text{cells}}\text{Dac}\tau_{\text{avg}}}, \quad (370)$$

where we have added a factor to include, as above, the idea that Hb expression levels at one cell depend on an average over N_{cells} nearby cells. Empirically, the mean expression level is well approximated by a Hill function,

$$\langle\text{Hb}\rangle = \frac{c^n}{c_{1/2}^n + c^n}, \quad (371)$$

where now we choose units where the maximum mean expression level is one, and the data are fit best by $n = 5$. Then we have

$$\frac{d\langle\text{Hb}\rangle}{d\ln c} = n\langle\text{Hb}\rangle(1 - \langle\text{Hb}\rangle), \quad (372)$$

and hence, after some algebra,

$$\langle(\delta\text{Hb})^2\rangle_{\text{BP}} = \beta\langle\text{Hb}\rangle^{2-1/n}(1 - \langle\text{Hb}\rangle)^{2+1/n}, \quad (373)$$

$$\beta = \frac{n^2}{N_{\text{cells}}\text{Dac}_{1/2}\tau_{\text{avg}}}. \quad (374)$$

If we have both the Berg–Purcell noise at the input to transcriptional control, and the Poisson noise at the output, then we expect the variances to add, so that

$$\langle(\delta\text{Hb})^2\rangle = \langle(\delta\text{Hb})^2\rangle_{\text{BP}} + \langle(\delta\text{Hb})^2\rangle_{\text{Poisson}} \quad (375)$$

$$= \beta\langle\text{Hb}\rangle^{2-1/n}(1 - \langle\text{Hb}\rangle)^{2+1/n} + \alpha\langle\text{Hb}\rangle. \quad (376)$$

In Figure 55 we see how this prediction compares with experiment. Since $n = 5$ is known from the input/output relation, we have to set the parameters α and β . At maximal mean expression, $\langle\text{Hb}\rangle = 1$ and Eq (376) predicts $\langle(\delta\text{Hb})^2\rangle = \alpha$, so we can read this parameter directly from the behavior at the right hand edge of the graph ($\alpha^2 \sim 0.05$). We have just one parameter β left to fit, but this will determine the height, shape and position of the peak in the noise level vs mean, so it is not at all guaranteed that we will get a reasonable fit. In fact the fit is very good, and we find $\beta \sim 0.5$. It is interesting that the dependence of the variance on the mean seems very sensitive, since if we let the Hill coefficient become large, even the best fit of Eq (376) systematically misses the data, as shown by the $n \rightarrow \infty$ curve in Fig 55. Other subtly different models also fail, as you can see in Problem 65 [careful with number!].




FIG. 55 Noise in Hunchback expression as a function of the mean expression level, from Tkačik et al (2008). This is a replotting of the data from Fig 54, compared with several models as described in the text. Error bars are standard deviations across multiple individual embryos.

Problem 64: Details of Hunchback noise. Discuss the meaning of the parameters α and β . Can you relate these to meaningful physical quantities? Do we have independent data to see if these numbers make sense?

Problem 65: Transcriptional bursting? The key point about noise in synthesis and degradation is that we expect the variance to be monotonic as a function of the mean (as in the Poisson model), and this is not what we see in Fig 55. An alternative model that could explain the peak of noise at intermediate expression levels is that the transcription site switches between active and inactive states, generating a “burst” of mRNA molecules while in the active state. You should be able to go back to our discussion of noise in binding and unbinding without diffusion [leading up to Eq (A322)], and build up the predictions of this model.

(a.) Suppose that switching into the active state occurs at a rate k_{on} , and the switch back to the inactive state occurs at a rate k_{off} . These rates must vary with the concentration of the input transcription factor, since it is only by switching between active and inactive states that the system can modulate the mean output. It seems plausible that the mean output is proportional to the probability of being in the active state. Are there any conditions under which this would not be true?

(b.) Show that if the mean output is proportional to the probability of being in the active state, then the random switching will contribute to the output variance a term

$$\langle(\delta\text{Hb})^2\rangle_{\text{burst}} = \langle\text{Hb}\rangle(1 - \langle\text{Hb}\rangle) \cdot \frac{\tau_c}{\tau_{\text{avg}}}, \quad (377)$$

where the correlation time $\tau_c = 1/(k_{\text{on}} + k_{\text{off}})$, the output is measured in units such that the maximal mean value is $\langle\text{Hb}\rangle = 1$, as above, and we assume that the averaging time is long compared with τ_c .

(c.) Switching into the active state is associated with transcription factor binding. In contrast, switching back to the inactive state doesn’t require any additional binding events. Thus it is plausible that the rate k_{off} is independent of the input concentration c . What is the dependence of k_{on} required to reproduce the mean input/output relation in Eq (371)? Is there a mechanistic interpretation of this dependence?

(d.) As an aside, can you give an alternative description based on the MWC model, as in our discussion of the bacterial rotary motor above? Notice that now you need to think about the kinetics of the transitions between the two states, not just the free energies. See also the Appendix A.4. This is deliberately open ended.

(e.) Combine your results in [b.] and [c.] to show that the analog of Eq (376) in this model is

$$\langle(\delta\text{Hb})^2\rangle = \langle(\delta\text{Hb})^2\rangle_{\text{burst}} + \langle(\delta\text{Hb})^2\rangle_{\text{Poisson}} \quad (378)$$

$$= \gamma\langle\text{Hb}\rangle(1 - \langle\text{Hb}\rangle)^2 + \alpha\langle\text{Hb}\rangle. \quad (379)$$

Give an expression for γ in terms of the original parameters of the model. Explain why the steepness of the Hill function (that is, the parameter n) doesn’t appear directly in determining the shape of the relation between variance and mean.

(f.) In Fig 55, we see the best fit of Eq (379) to the data, which is not very good. Without doing a fit, you should be able to show that the model predicts a relation between the point at which the noise is maximal, and the height of this maximum. Show that this is inconsistent with the data.

To summarize, we can now observe directly the noise in gene expression. While one could emphasize that these fluctuations are, under some conditions, quite large, it seems more surprising that there are conditions where they are quite small. Cells can set the output of their genetic control machinery with a precision of $\sim 10\%$ or better, thus doing much more than switching genes on

and off—intermediate levels of expression are meaningful. This means, in particular, that we have made measurements with an accuracy of better than 10%, and this isn’t always easy to do. More fundamentally, the precision with which cells can control expression levels is not far from the limits set by the random arrival of the relevant signaling molecules (transcription factors) at their targets. Of course, we could imagine cells which use more copies of all the transcription factors, and thus could achieve greater precision—or be sloppier, and reach the same precision—but this doesn’t seem to be what happens. I don’t think we understand why evolution has pushed cells into this particular corner.

So far we have discussed noise as a small fluctuation around the mean. It is also possible that, in the same way that thermal noise can result in a nonzero rate for chemical reactions, noise in chemical kinetics can generate spontaneous switching among otherwise stable states. Much has been written about this. I am less certain that we really understand any particular system. There is, however, some elegant physics here, so I would like to come back and discuss this.

The following two problems are concerned with a newly discovered bacterium that responds to a chemical signal by emitting light. The bacteria are roughly spherical, with diameter $d \sim 2 \mu\text{m}$, and hence are clearly visible under the microscope. The chemical signal is shown to be a small protein, presumably secreted by other bacteria; the protein diffuses through the extracellular medium with a diffusion constant $D \sim 10 \mu\text{m}^2/\text{s}$. Very careful experiments establish that each individual bacterium either emits light at full intensity or is essentially dark, and that changing the concentration c of the signaling protein changes the probability of being in the two states. Larger values of c correspond to higher probabilities of being in the light emitting state, so that $p_{\text{light}}(c)$ is monotonically increasing.

Problem 66: Extreme sensitivity, but slowly. There is a specific concentration $c = c_{1/2}$ of the signaling protein such that $p_{\text{light}}(c_{1/2}) = p_{\text{dark}}(c_{1/2}) = 0.5$. When poised at $c = c_{1/2}$ the system switches back and forth between the two states spontaneously at a rate of $\sim 1/\text{hour}$. Remarkably, a change in c by just 10% is sufficient to shift the probabilities from $p_{\text{light}} = 0.5$ to $p_{\text{light}} = 0.9$ or $p_{\text{light}} = 0.1$ when the concentration is increased or decreased, respectively.

(a.) After some confusion in early experiments, it is found that everything said above is true, but the half-maximal concentration $c_{1/2} = 10^{-12} \text{M}$. Is this possible? Justify your answer clearly and quantitatively.

(b.) One group proposes that this extreme sensitivity is not at all surprising, since after all proteins can bind to other proteins with dissociation constants as small as $K_D \sim 10^{-15} \text{M}$. Does this observation of very tight binding have anything to do with the physical limits on sensitivity? Why or why not?

(c.) Another group notes that 10^{-12}M corresponds to $\sim 10^{-3}$ molecules in the volume of the bacterium. They argue that this provides evidence for homeopathy, in which drugs are claimed to retain their effectiveness at extreme dilution, perhaps even to the point where the doses contain less than one molecule on average. Can you resolve their confusion?

Problem 67: How simple can it be? Further studies of this new light emitting bacterium aim at identifying the molecules involved. The first such experiment shows that if you block protein

synthesis, the system cannot switch between the dark and light states, indicating that the switch involves a change in gene expression rather than (for example) a change in phosphorylation or methylation states of existing proteins as in chemotaxis. A systematic search which knocks out individual genes, looking for effects on the behavior, finds only one gene that codes for a DNA-binding protein. When this gene is knocked out, all bacteria are permanently dark. More detailed experiments show that these bacteria not only are dark, they actually are not expressing the proteins required for generating light.

(a.) Draw the simplest schematic model suggested by these results. Be sure that your model explains why there are two relatively stable states (light and dark) rather than a continuum of intermediates, and that your model is consistent with the knock out experiments.

(b.) Assume that the signaling protein binds to some receptor on the surface of the cell and that this triggers a cascade of biochemical events. For simplicity you can imagine that the output of this cascade is some molecule, the concentration of which is proportional to the average occupancy of the receptors over some window of time. Explain how this molecule can couple to your model in [a] to influence the probability of the cell being in the dark or light states.

(c.) Formalize your models from [a] and [b] by writing differential equations for the concentrations of all the relevant species. Show how these equations imply the existence of discrete light/dark states. Can you see directly from the equations why changing the receptor occupancy will shift the balance between these states? It might be hard to explain the behavior near the midpoint ($c = c_{1/2}$), but it should be possible to explain the dominance of the dark state as $c \rightarrow 0$ and the light state as $c \rightarrow \infty$.

(d.) Describe qualitatively all the sources of noise that could enter your model. Do you have any guidance from experiment about which sources are dominant?

(e.) Consider the point where $c = c_{1/2}$. Explain qualitatively what features of your model are responsible for determining the ~ 1 hour time scale for jumping back and forth between the light and dark states.

(f.) See how far you can go in turning your remarks in [e] into an honest calculation!

There are several messages which I hoped to convey in this section. First, bacterial chemotaxis provides us with an example of chemical sensing which is interesting, not just in itself but as an example of a vastly more general phenomenon. Importantly, experiments on chemotaxis set a quantitative standard that should be emulated in the exploration of other chemical signaling systems, from the embryo to the brain. Second, as explained in Appendix A.6, the intuitive argument of Berg and Purcell can be made rigorous. What they identified is a limit to chemical signaling which is very much analogous to the photon shot noise limit in vision or imaging more generally. While molecules do many complicated things, they have to reach their targets in order to do them, and this is a random process, so this randomness sets a limit to the precision of almost everything that cells do.⁴⁷ Finally, real cells operate close to this limit, not just in

specialized tasks such as chemotaxis but in the everyday business of regulating gene expression. While other noise sources are clearly present, the “noise floor” that results from the Berg–Purcell limit never seems far away, and in some cases cells may push all the way to the point where this is the dominant noise source.

The study of chemotaxis has a long history. From a biologist’s point of view, the modern era starts when Adler (1965, 1969) demonstrates, using mutants, that chemosensing is independent of metabolism. From a physicist’s point of view, the modern era starts when Berg builds his tracking microscope and observes, quantitatively, the paths of individual bacteria (Berg 1971, Berg & Brown 1972). The experiments which demonstrated the temporal character of the computations involved in chemotaxis were done by Macnab & Koshland (1972) and by Brown & Berg (1974). A nice discussion of how these temporal comparisons translate into mobility up the gradient of attractive chemical is given by Schnitzer et al (1990).

Adler 1965: Chemotaxis in *Escherichia coli*. *Cold Spring Harbor Symp Quant Biol* **30**, 289–292 (1965).

Adler 1969: Chemoreceptors in bacteria. J Adler, *Science* **166**, 1588–1597 (1969).

Berg 1971: How to track bacteria. HC Berg, *Rev Sci Instrum* **42**, 868–871 (1971).

Berg & Brown 1972: Chemotaxis in *Escherichia coli* analyzed by three-dimensional tracking. *Nature* **239**, 500–504 (1972).

Brown & Berg 1974: Temporal stimulation of chemotaxis in *Escherichia coli* *Proc Nat’l Acad Sci (USA)* **71**, 1388–1392 (1974).

Macnab & Koshland 1972: R Macnab & DE Koshland, The gradient-sensing mechanism in chemotaxis. *Proc Nat’l Acad Sci (USA)* **69**, 2509–2512 (1972).

Schnitzer et al 1990: Strategies for chemotaxis. M Schnitzer, SM Block, HC Berg & EM Purcell, *Symp Soc Gen Microbiol* **46**, 15–34 (1990).

For fluid mechanics in general, see Landau and Lifshitz (1987). The fact that bacteria live at low Reynolds number, and that this must matter for their lifetsyle, was surely was known to many people, for many years. But Berg’s experiments on *E coli* provided a stimulus to think about this, and it resulted in a beautiful exposition by Purcell (1977), which has been hugely influential. The appreciation that self-propulsion at low Reynolds number has a gauge theory

tics for the photon counts, and this raises the question of whether we could generate comparable noise reductions in chemical processes. I think the answer is yes—for example, one could transport molecules to their targets by an active process that is more orderly than diffusion—but this seems enormously costly, as first emphasized by Berg and Purcell themselves. It is, however, worth thinking about. More subtly, some chemical reactions involve enormous numbers of steps, so that the fractional variance in the time required for completion of the reaction by one molecule becomes very small, as in the discussion of rhodopsin shutoff in Section I.C. Indeed, transcription itself can be seen as an example, where it is possible for the time required to synthesize a single mRNA molecule—once transcription has been initiated—to be nearly deterministic, so that this process does not contribute a significant amount of noise.

⁴⁷ It is possible to produce light that does not obey Poisson statis-

description is due to Shapere & Wilczek (1987). The dramatic discovery that bacteria swim by rotating their flagella was made by Berg & Anderson (1973), and then Silverman & Simon (1974) succeeded in tethering cells by their flagella to see the rotation of the cell body.

Berg & Anderson 1973: Bacteria swim by rotating their flagellar filaments. *Nature* **245**, 380–382 (1973).

Landau & Lifshitz 1987: *Fluid Mechanics*. LD Landau & EM Lifshitz (Pergamon, Oxford, 1987).

Purcell 1977: EM Purcell, Life at low Reynolds' number, *Am J Phys* **45**, 3–11 (1977).

Shapere & Wilczek 1987: Self-propulsion at low Reynolds number. *Phys Rev Lett* **58**, 2051–2054 (1987).

Silverman & Simon 1974: Flagellar rotation and the mechanism of bacterial motility. M Silverman & M Simon, *Nature* **249**, 73–74 (1974).

Should add some references about rotation of the mitochondrial ATPase, and more recent work on flagellar motor

:

The classic, intuitive account of the physical limits to chemical sensing is by Berg and Purcell (1977). [Do we want to dig into the papers that they reference, in relation to sensitivity?] Measurements on the impulse response of the system were reported by Block et al (1982), and these experiments, along with Segall et al (1986) provide a more compelling demonstration that bacterium is sensitive to single molecular events. Another interesting paper from this period is Block et al (1983) [should tell the story about the Appendix as an example of models/theories in biology]. The idea of deriving the impulse response as the solution to an optimization problem, in the spirit of the Berg–Purcell discussion but more rigorously, has been explored by several groups: Strong et al (1998), Andrews et al (2006), and most recently Celani & Vergassola (2010), who introduced a novel game theoretic approach [check other refs].

Andrews et al 2006: Optimal noise filtering in the chemotactic response of *Escherichia coli*. BW Andrews, T–M Yi & PA Iglesias, *PLoS Comp Bio* **2**, e154 (2006).

Berg & Purcell 1977: Physics of chemoreception. HC Berg & EM Purcell, *Biophys J* **20**, 193–219 (1977).

Block et al 1982: Impulse responses in bacterial chemotaxis. SM Block, JE Segall & HC Berg, *Cell* **31**, 215–226 (1982).

Block et al 1983: Adaptation kinetics in bacterial chemotaxis. SM Block, JE Segall & HC Berg, *J Bacteriol* **154**, 312–323 (1983).

Celani & Vergassola 2010: Bacterial strategies for chemotaxis. A Celani & M Vergassola, *Proc Nat'l Acad Sci (USA)* **107**, 1391–1396 (2010).

Segall et al 1986: Temporal comparisons in bacterial chemotaxis. JE Segall, SM Block & HC Berg, *Proc Nat'l Acad Sci (USA)* **83**, 8987–8991 (1986).

Strong et al 1998: Adaptation and optimal chemotactic strategy in *E. coli*. SP Strong, B Freedman, W Bialek & R Koberle, *Phys Rev E* **57**, 5604–5617 (1998).

The experiments on the response of the flagellar motor to the CheY~P concentration are by Cluzel et al (2000). For measurements on the diffusion constant of proteins in *E. coli* see Elowitz et al (1999), and for observations on the structure of the motor in relation to its regulation by CheY~P, see Thomas et al (1999). The model in Fig 49 is based on give original refs for MWC–style description of rotation. Give refs to models at the front end of the transduction scheme, depending on what gets said in the text!

Cluzel et al 2000: An ultrasensitive bacterial motor revealed by monitoring signaling proteins in single cells. P Cluzel, M Surette & S Leibler, *Science* **287**, 1652–1655 (2000).

Elowitz et al 1999: Protein mobility in the cytoplasm of *Escherichia coli*. MB Elowitz, MG Surette, P–E Wolf, JB Stock & S Leibler, *J Bacteriol* **181**, 197–203 (1999).

Thomas et al 1999: Rotational symmetry of the C ring and a mechanism for the flagellar rotary motor. DR Thomas, DG Morgan & DJ DeRoiser, *Proc Nat'l Acad Sci (USA)* **96**, 10134–10139 (1999).

This seems to be the first place where GFP–based methods have come up, so need to give a guide of the literature here!

:

In thinking about transcriptional regulation, it is useful to review some basic facts about molecular biology, for which the classic reference is Watson's *Molecular Biology of the Gene*. This has been through many editions, and at times flirted with being more of an encyclopedia than a textbook. I'll reference the current edition here, which seems a bit more compact than some of the intermediate editions, but I also encourage you to look back at earlier editions, written by Watson alone. A beautiful account of gene regulation, using the bacteriophage λ as an example, was given by Ptashne (1986), which has also evolved with time (Ptashne 1992); see also Ptashne (2001).

Ptashne 1986: *A Genetic Switch: Gene Control and Phage λ* . M Ptashne (Cell Press, Cambridge MA, 1986).

Ptashne 1992: *A Genetic Switch, Second Edition: Phage λ and Higher Organisms*. M Ptashne (Cell Press, Cambridge MA, 1992).

Ptashne 2001: *Genes and Signals*. M Ptashne (Cold Spring Harbor Laboratory Press, New York, 2001).

Watson et al 2008: *Molecular Biology of the Gene, Sixth Edition*. JD Watson, TA Baker, SP Bell, A Gann, M Levine & R Losick (Benjamin Cummings, 2008).

In order to make our discussion quantitative, we need to know the absolute concentration at which transcription factors act. Ptashne's books give some discussion of this, although the estimates were a bit indirect. Several groups have made measurements on the binding of transcription factors to DNA, trying to measure the concentration at which binding sites are half occupied; sometimes this is done by direct physical–chemical methods *in vitro*, and sometimes by less direct methods *in vivo*. Examples include Oehler et al (1994), Ma et al (1996), Pedone et al (1996), Burz et al (1998), and Winston et al (1999). A modern version of the *in vitro* binding experiment examines the molecules one at a time, as in the work by Wang et al (2009).

Burz et al 1998: Cooperative DNA binding by Bicoid provides a mechanism for threshold dependent gene activation in the *Drosophila* embryo. DS Burz, R Pivera–Pomar, H Jackle & SD Hanes, *EMBO J* **17**, 5998–6009 (1998).

Ma et al 1996: The *Drosophila* morphogenetic protein Bicoid binds DNA cooperatively. X Ma, D Yuan, K Diepold, T Scarborough, & J Ma, *Development* **122**, 1195–1206 (1996).

Oehler et al 1994: Quality and position of the three lac operators of *E. coli* define efficiency of repression. S Oehler, M Amouyal, P Kolkhof, B von Wilcken–Bergmann & B Müller–Hill, *EMBO J* **13**, 3348–3355 (1994).

Pedone et al 1996: The single Cys2–His2 zinc finger domain of the GAGA protein flanked by basic residues is sufficient for high–affinity specific DNA binding. PV Pedone, R Ghirlando, GM Clore, AM Gronenborn, G Felsenfeld & JG Omichinski, *Proc Nat'l Acad Sci (USA)* **93**, 2822–2826 (1996).

Wang et al 2009: Quantitative transcription factor binding kinetics at the single molecule level. Y Wang, L Guo, I Golding, EC Cox, NP Ong, *Biophys J* **96**, 609–620 (2009).

Winston et al 1999: Characterization of the DNA binding properties of the bHLH domain of Deadpan to single and tandem sites. RL Winston, DP Millar, JM Gottesfeld, Gottesfeld & SB Kent. *Biochemistry* **38**, 5138–5146 (1999).

An important development in the field has been the construction of fusion proteins, combining transcription factors with fluorescent proteins, and the re-insertion of these fusions into the genome. For more about these techniques in general, see the references at the end of Section II.B. When cells divide, their contents are partitioned, and one can observe the noise from the finite number of molecules being assigned at random to one of the two daughter cells. Rosenfeld et al (2005), and more recently Teng et al (2010) has shown how this can be used to make very precise estimates of the number of copies of the protein in the mother cell, and thus providing a calibration that converts fluorescence intensity back into copy number. Gregor et al (2007a) discuss a case where it was possible to test in detail that the fusion construct replaces the function of the original transcription factor, quantitatively, and in the next paper they exploit this construct to analyze the noise in one step of transcriptional regulation (see below), as well as making estimates of absolute concentration by comparing the fluorescence intensity to a purified standard (Gregor et al 2007b).

Gregor et al 2007a: Stability and nuclear dynamics of the Bicoid morphogen gradient. T Gregor, EF Wieschaus, AP McGregor, W Bialek & DW Tank, *Cell* **130**, 141–152 (2007).

Gregor et al 2007b: Probing the limits to positional information. T Gregor, DW Tank, EF Wieschaus & W Bialek, *Cell* **130**, 153–164 (2007).

Rosenfeld et al 2005: Gene regulation at the single cell level. N Rosenfeld, JW Young, U Alon, PS Swain & MB Elowitz, *Science* **307**, 1962–1965 (2005).

Teng et al 2010: Measurement of the copy number of the master quorum-sensing regulator of a bacterial cell. S–W Teng, Y Wang, KC Tu, T Long, P Mehta, NS Wingreen, BL Bassler & NP Ong, *Biophys J* **98**, 2024–2031 (2010).

In contrast to bacteria, many eukaryotic cells are large enough, or move slowly enough, that they can get a reliable signal by measuring gradients across the length of their body; for a discussion of the limits to these measurements and some of the relevant experiments, see Endres & Wingreen (2009a,b). [Need to digest data on chemotaxis in bigger cells Find general reference on axon guidance, growth cones etc..](#) The measurements on extreme precision of axon guidance were reported by Rosoff et al (2004).

Endres & Wingreen 2009a: Accuracy of direct gradient sensing by single cells. RG Endres & NS Wingreen, *Proc Natl Aca Sci (USA)* **105**, 15749–15754 (2008).

Endres & Wingreen 2009b: Accuracy of direct gradient sensing by cell-surface receptors. RG Endres & NS Wingreen, *Prog Biophys Mol Biol* **100**, 33–39 (2009).

Gregor et al 2010: The onset of collective behavior in social amoebae. T Gregor, K Fujimoto, N Masaki & S Sawai, *Science* **328**, 1021–1025 (2010).

Rosoff et al 2004: A new chemotaxis assay shows the extreme sensitivity of axons to molecular gradients. WJ Rosoff, JS Urbach, MA Esrick, RG McAllister, LJ Richards & GJ Goodhill, *Nature Neurosci* **7**, 678–682 (2004).

Song et al 2006: *Dictyostelium discoideum* chemotaxis: Threshold for directed motion. L Song, SM Nadkarnia, HU Bödeker, C Beta, A Bae, C Franck, W-J Rappel, WF Loomis & E Bodenschatz, *Eur J Cell Bio* **85**, 981–989 (2006).

It is only in the last decade that it has been possible to make direct measurements of the noise in gene expression, and even more recently that it has been possible to focus on noise in the control process itself. The initial experiment separating intrinsic from extrinsic noise sources using the two color plasmid was by Elowitz et

al (2002), which touched off a series of experiments on both bacterial (Ozbudak et al 2002, Pedraza & van Oudenaarden 2005) and eukaryotic systems (Blake et al 2003, Raser & O’Shea 2004). The experiments on noise in the Bcd/Hb system are by Gregor et al (see above). A review of methods for measuring Bcd concentration profiles is given by Morrison et al (2011), and in particular they discuss the comparison of live GFP-based imaging with antibody staining methods in fixed samples. A more detailed analysis of the data on Bcd/Hb noise is given by Tkačik et al (2008), which also provides a broader context on the role of different noise sources in the control of gene expression. Models based on transcriptional bursting are inspired by the direct observation of these bursts in *E coli* by Golding et al (2005). It is worth thinking about whether the observed bursts necessarily result from the kinetics of switching between states of the transcriptional apparatus, or could be traced to the binding and unbinding of transcription factors.

Blake et al 2003: Noise in eukaryotic gene expression. WJ Blake, M Kaern, CR Cantor & JJ Collins, *Nature* **422**, 633–637 (2003).

Elowitz et al 2002: Stochastic gene expression in a single cell. MB Elowitz, AJ Levine, ED Siggia & PD Swain, *Science* **297**, 1183–1186 (2002).

Golding et al 2005: Real-time kinetics of gene activity in individual bacteria. I Golding, J Paulsson, SM Zawilski & EC Cox, *Cell* **123**, 1025–1036 (2005).

Morrison et al 2011: Quantifying the Bicoid morphogen gradient in living embryos. AH Morrisson, M Scheeler, J Dubuis & T Gregor, in *Imaging in Developmental Biology: A Laboratory Manual*, J Sharpe & R Wong, eds (Cold Spring Harbor Press, Woodbury NY, 2011); arXiv.org:1003.5572 [q-bio.QM] (2010).

Ozbudak et al 2002: Regulation of noise in the expression of a single gene. E Ozbudak, M Thattai, I Kurtser, AD Grossman & A van Oudenaarden, *Nature Gen* **31**, 69–73 (2002).

Pedraza & van Oudenaarden 2005: Noise propagation in gene networks. J Pedraza & A van Oudenaarden, *Science* **307**, 1965–1969 (2005).

Raser & O’Shea 2004: Control of stochasticity in eukaryotic gene expression. JM Raser & EK O’Shea, *Science* **304**, 1811–1814 (2004).

Tkačik et al 2008: The role of input noise in transcriptional regulation. G Tkačik, T Gregor & W Bialek, *PLoS One* **3**, e2774 (2008).

[Will need to add some references about bistability, noise induced switching, and maybe path integral methods for noise .. depends on what gets said in the text.](#)

C. More about noise in perception

We have already said a bit about noise in visual perception, in the case where perception amounts to counting photons. But this is just one corner of our perceptual experience, and we’d like to know if some of the same principles are relevant outside of this limit. In this section we will look at a few instances, sampled from different organisms and different sensory modalities. I think one of the important ideas here is that considerations of noise—and processing strategies for reaching reliable

conclusions in the presence of noise, perhaps even optimizing performance—cut across these many different systems, which often are the subjects of quite isolated literatures.

It has been known for some time that bats navigate by generating ultrasonic calls and listening for the echoes, forming an image of their world much as in modern sonar. To get a feeling for the precision of this behavior, there is a simple, qualitative experiment that is best explained with a certain amount of (literal) hand waving [ask Jim Simmons for original reference]. Some bats will happily eat mealworms if you toss them into the air. Before tossing them, however, you can dip them into a little bit of flour. To eat the worm, the bat must “see” it, and then maneuver its own body into position, finally sweeping the worm up in its wing and bringing it to its mouth. But if the worm has been dusted with flour, this will leave a mark on the wing. Now repeat the experiment, many times, with same bat (but, of course, different worms). If you look at the bat’s wing, you might expect to see many spots of flour, but in fact all the spots are on top of one another. This suggests that the entire process—not just identifying the location of the worm in the air, but the acrobatic movements required to scoop it up—have a precision of roughly one centimeter. In echolocation, position estimates are based on the time delays of the echoes, and with a sound speed of ~ 340 m/s, this corresponds to a timing precision of $\delta t \sim 30 \mu\text{s}$. This rough estimate already is interesting, although maybe not too shocking since we can detect a few microseconds of difference in the arrival times of sounds between our two ears, and this is how we can localize the source of low frequency sounds. Barn owls do even better, detecting $\delta\tau \sim 1 \mu\text{sec}$ between their ears.

As an aside, it was Rayleigh who understood that our




FIG. 56 A schematic of the ‘Y’ apparatus for testing echo timing discrimination performance in bats, from Simmons et al (1990).




FIG. 57 Performance of four different bats at echo jitter discrimination, from Simmons et al (1990). Echoes can be returned with no phase shift (circles), or with a phase shift of π (squares); errors for the phase shifted echoes are measured downward. We see that the phase shift itself is detectable with almost no errors, that there is confusion around $\delta\tau \sim 35 \mu\text{s}$, and that this “confusion peak” shifts and splits with the introduction of a phase shift.

brains need to use different cues for localization in different frequency ranges, just because of the physics of sound waves. At high frequencies (short wavelengths) our head casts an acoustic shadow, and there is a difference in intensity between our ears—the sound comes from the side that gets the louder signal. But at low frequencies, the wavelength is comparable to or larger than the size of our head, and there is no shadow. There is, however, a time or phase difference, but this is small. To demonstrate our sensitivity to these small time differences directly, he sat Lady Rayleigh in the gazebo behind their home, and arranged for tubes of slightly different length to lead from a sound source to her two ears. A fabulous image.

Problem 68: Time differences and binaural hearing.

Show that when a sound source is far away, the difference in propagation time to your two ears is independent of distance to the source. What does determine this time difference? For your own head, what is the time difference for a source at an angle of $\sim 10^\circ$ to the right of the direction your nose is pointing?

To be more quantitative, one would like to get the bats to report more directly on their estimates of echo delay, as in Fig 56. In one class of experiments, bats stand at the base of a Y with loudspeakers on the two arms. Their ultrasonic calls are monitored by microphones and returned through the loudspeakers with programmable delays. In

a typical experiment, the ‘artificial echoes’ produced by one side of the Y are at a fixed delay τ , while the other side alternately produces delays of $\tau \pm \delta\tau$. The bat is trained to take a step toward the side which alternates, and the question is how small we can make $\delta\tau$ and still have the bat make reliable decisions. Early experiments suggested that delays differences of $\delta\tau \sim 1 \mu\text{sec}$ were detectable, and perhaps more surprisingly that delays of $\sim 35 \mu\text{sec}$ were less detectable, as shown in Fig 57. The latter result might make sense if the bat were trying to measure delays by matching the detailed waveforms of the call and echo, since these sounds have most of their power at frequencies near $f \sim 1/(35 \mu\text{sec})$ —the bat can be confused by delay differences which correspond to an integer number of periods in the acoustic waveform, and one can even see the $n = 2$ ‘confusion resonance’ if one is careful. One can also introduce a phase shift into the artificial echo, and this shifts the confusion peak as expected.

Let’s think about this more formally. Suppose that we

are expecting a sound (or any signal) that has a time dependence $s_0(t)$, but we don’t know when it will arrive, so what we actually observe will be $s_0(t - \tau)$ embedded in some background of noise. That is,

$$s(t) = s_0(t - \tau) + \eta(t), \quad (380)$$

where $\eta(t)$ is the noise. Let’s assume, for simplicity, that the noise is white, with some spectral density \mathcal{N} . Then, as explained in Appendix B, the probability density for the function $s(t)$ becomes

$$P[s(t)|\tau] = \frac{1}{Z} \exp \left[-\frac{1}{2\mathcal{N}} \int dt \left| s(t) - s_0(t - \tau) \right|^2 \right], \quad (381)$$

where Z is a normalization constant and the notation reminds us that this is the distribution if we know the delay τ . If instead the delay is $\tau + \delta\tau$,

$$P[s(t)|\tau + \delta\tau] = \frac{1}{Z} \exp \left[-\frac{1}{2\mathcal{N}} \int dt \left| s(t) - s_0(t - \tau - \delta\tau) \right|^2 \right]. \quad (382)$$

As in our previous discussions of discrimination between two alternatives [give specific pointer], when we are faced with a particular signal $s(t)$ and have to decide whether the delay was τ or $\tau + \delta\tau$, the relevant quantity is the (log) likelihood ratio:

$$\lambda[s(t)] \equiv \ln \left(\frac{P[s(t)|\tau + \delta\tau]}{P[s(t)|\tau]} \right) \quad (383)$$

$$= -\frac{1}{2\mathcal{N}} \int dt \left| s(t) - s_0(t - \tau - \delta\tau) \right|^2 + \frac{1}{2\mathcal{N}} \int dt \left| s(t) - s_0(t - \tau) \right|^2 \quad (384)$$

$$= \frac{1}{\mathcal{N}} \int dt s(t) [s_0(t - \tau - \delta\tau) - s_0(t - \tau)]. \quad (385)$$

If the delay really is τ , then

$$\langle \lambda[s(t)] \rangle_\tau \equiv \left\langle \frac{1}{\mathcal{N}} \int dt s(t) [s_0(t - \tau - \delta\tau) - s_0(t - \tau)] \right\rangle_\tau \quad (386)$$

$$= \left\langle \frac{1}{\mathcal{N}} \int dt [s_0(t - \tau) + \eta(t)] [s_0(t - \tau - \delta\tau) - s_0(t - \tau)] \right\rangle_\tau \quad (387)$$

$$= \frac{1}{\mathcal{N}} \int dt s_0(t - \tau) [s_0(t - \tau - \delta\tau) - s_0(t - \tau)] \quad (388)$$

$$= \frac{1}{\mathcal{N}} [C(\delta\tau) - C(0)], \quad (389)$$

where

$$C(t) = \int dt' s_0(t') s_0(t' - t) \quad (390)$$

is the autocorrelation function of the expected signal.

Similar calculations yield

$$\langle \lambda[s(t)] \rangle_{\tau + \delta\tau} = \frac{1}{\mathcal{N}} [C(0) - C(\delta\tau)], \quad (391)$$

$$\langle (\delta\lambda[s(t)])^2 \rangle_\tau = \langle (\delta\lambda[s(t)])^2 \rangle_{\tau + \delta\tau} \quad (392)$$

$$= \frac{2}{\mathcal{N}} [C(0) - C(\delta\tau)]. \quad (393)$$

It should also be clear that $\lambda[s(t)]$ is a Gaussian random variable (inherited from the Gaussian statistics of the noise η), so these few moments provide a complete description of the problem of discriminating between delays τ and $\tau + \delta\tau$. The end result is that the discrimination problem is exactly that of a single Gaussian variable (λ), with signal-to-noise ratio

$$SNR = \frac{(\langle \lambda[s(t)] \rangle_{\tau+\delta\tau} - \langle \lambda[s(t)] \rangle_{\tau})^2}{\langle (\delta\lambda[s(t)])^2 \rangle} = \frac{2}{\mathcal{N}} [C(0) - C(\delta\tau)]. \quad (394)$$

Thus we see that the SNR is large as soon as the jitter $\delta\tau$ is big enough to break the correlations in the waveform, and conversely that the SNR falls if shifting by $\delta\tau$ brings the waveform back into correlation with itself, as will happen for an approximately periodic signal such as the echolocation pulse.

Problem 69: Details of the SNR for detecting jitter in echolocation. Fill in the details leading to Eq (394).

(a.) How does this result change if the discrimination involves not just a time shift $\delta\tau$ but also a sign flip or π phase shift?

(b.) Recall the relationship between error probability and SNR [point back to photon counting discussion]. Is it practical to try and estimate the correlation function $C(\tau)$ by measuring the error probability as a function of $\delta\tau$? What if you also have access to experiments with a sign flip, as in (a.)? If you have errors in the measurement of the error probability, how do these propagate back to estimates of the underlying $C(\tau)$?

(c.) Compare your results in (b.) with the construction of “compound jitter discrimination curves” by Simmons et al (1990). Could you suggest improvements in their data analysis methods?

This argument about discriminability assumes that the bat’s brain actually can compute using the entire acoustic waveform $s(t)$, rather some more limited features; in this sense we are describing the best that the bat could possibly do. It is interesting that such a calculation predicts confusion at delays where the autocorrelation function of the bat’s call has a peak, and that such confusions are observed. On the other hand, this calculation seems hopelessly optimistic: “access to the acoustic waveform” means, in particular, access to features that are varying on the microsecond timescale. If we record the activity of single neurons emerging from the ear as they respond to pure tones, then we can see the action potentials “phase lock” to the tone, but this effect is significant only up to some maximum frequency. Beyond this high frequency cutoff, the overall rate of spikes increases with the intensity of the tone, but the timing of the spikes seems unrelated to the details of the acoustic waveform. Although there is controversy about the precise value of the cutoff frequency for phase locking, there seems to be no hint in

the literature that it could be as high as 30 kHz. Taking all this at face value, it seems implausible that the auditory nerve actually transmits to the brain anything like a complete replica of the echo waveforms.

There is a second problem with this seemingly simple calculation. If we expand the SNR for small $\delta\tau$, we have

$$SNR = \frac{2}{\mathcal{N}} [C(0) - C(\delta\tau)] \approx \frac{C(0)}{\mathcal{N}} \cdot \left[\frac{C''(0)}{C(0)} \right] (\delta\tau)^2. \quad (395)$$

We expect that the term in brackets, which has the units of $1/(\text{time})^2$, is determined by the time scale on which the echolocation pulse is varying, something like $\sim 35 \mu\text{sec}$. On the other hand, the first term, $C(0)/\mathcal{N}$ measures how loud the echo is relative to the background noise, and is dimensionless. We recall that in acoustics it is conventional to measure in decibels, where 10 dB represents a factor of ten difference in acoustic power or energy. A typical quiet conversation produces sounds ~ 30 dB above our threshold of hearing and hence above the limiting internal noise sources in the ear, whatever these may be. The bat’s echolocation pulses are enormously loud, and although the echoes may be weak, it still is plausible that (at least in the laboratory setting) they are ~ 60 dB above the background noise. This means that our calculation predicts a signal-to-noise ratio of one when the differences in delay $\delta\tau$ are measured in tens of *nanoseconds*, not microseconds. I think this was viewed as so obviously absurd that it was grounds for throwing out the whole idea that the bat uses detailed waveform information, even without reference to data on what the auditory nerve can encode.

In an absolutely stunning development, however, Simmons and colleagues went back to their experiments, produced delays in the appropriate range—convincing yourself that you have control of acoustic and electronic delays with nanosecond precision is not so simple—and found that the bats could do what they should be able to do as ideal detectors: they detect 10 *nanosecond* differences in echo delay, as shown in Fig 58. Further, they added noise in the background of the echoes and showed that performance of the bats tracked the ideal performance over a range of noise levels. This is a wonderful example with which to start this section of our discussion, since we have absolutely no idea how the bat manages this amazing feat of signal processing.

The problem of echo delay discrimination has just enough structure to emphasize an important point: when we make perceptual decisions, we are not identifying signals, we are identifying the distribution out of which these signals have been drawn. This becomes even more important as we move toward more complex tasks, where the randomness is intrinsic to the ‘signal’ rather than just a result of added noise. As an example, a single spoken word can generate a wide variety of sounds, all the more varied when embedded in a sentence. Identi-




FIG. 58 Bat echo discrimination performance at very small delays, from Simmons et al (1990). **Should add something about dependence on background noise level.]**

fying the word really means saying that the particular sound we have heard comes from this distribution and not another. Importantly, probability distributions can overlap, and hence there are limits on the reliability of discrimination.

Some years ago, Barlow and colleagues launched an effort to use these ideas of discrimination among distributions to study progressively more complex aspects of visual perception, in some cases reaching into the psychology literature for examples of gestalt phenomena—where our perception is of the whole rather than its parts. One such example is the recognition of symmetry in otherwise random patterns. Suppose that we want to make a random texture pattern. One way to do this is to draw the contrast $C(\mathbf{x})$ at each point \mathbf{x} in the image from some simple probability distribution that we can write down. An example is to make a Gaussian random texture, which corresponds to

$$P[C(\mathbf{x})] \propto \exp \left[-\frac{1}{2} \int d^2x \int d^2x' C(\mathbf{x}) K(\mathbf{x} - \mathbf{x}') C(\mathbf{x}') \right], \quad (396)$$

where γ measures the strength of the tendency toward symmetry. Clearly as $\gamma \rightarrow \infty$ we have an exactly symmetric pattern, quenching half of the degrees of freedom in the original random texture. On the other hand, as $\gamma \rightarrow 0$, the weakly symmetric textures drawn from P_γ become almost indistinguishable from a pure random tex-

ture ($\gamma = 0$). Given images of a certain size, and a known kernel K , there is a limit to the smallest value of γ that can be distinguished reliably from zero, and we can compare this statistical limit to the performance of human observers. This is more or less what Barlow did, although he used blurred random dots rather than the Gaussian

Problem 70: Texture discrimination. Show that Eq (396) can be rewritten as

$$P[C(\mathbf{x})] \propto \exp \left[-\frac{1}{2} \int \frac{d^2k}{(2\pi)^2} \frac{|\tilde{C}(\mathbf{k})|^2}{S_C(\mathbf{k})} \right], \quad (397)$$

where $S_C(\mathbf{k})$ is the (now two dimensional) power spectrum, connected as usual to the correlation function

$$\langle C(\mathbf{x})C(\mathbf{x}') \rangle = \int \frac{d^2k}{(2\pi)^2} S_C(\mathbf{k}) e^{i\mathbf{k} \cdot (\mathbf{x} - \mathbf{x}')}. \quad (398)$$

Suppose that you have the task of discrimination between images drawn from distributions characterized by two different power spectra, $S_C(\mathbf{k})$ and $S_C(\mathbf{k}) + \Delta S_C(\mathbf{k})$. Show that, assuming one has access to a large area of the image, the discrimination problem for small $\Delta S_C(\mathbf{k})$ is again like the discrimination of a single Gaussian variable. Explain what role is played by the assumption of a “large area,” and what defines large in this context. How does the signal-to-noise ratio for discrimination depend on area?

The statement that texture has symmetry across an axis is that for each point \mathbf{x} we can find the corresponding reflected point $\hat{\mathbf{R}} \cdot \mathbf{x}$, and that the contrasts at these two points are very similar; this should be true for every point. This can be accomplished by choosing

$$P_\gamma[C(\mathbf{x})] \propto \exp \left[-\frac{1}{2} \int d^2x \int d^2x' C(\mathbf{x}) K(\mathbf{x} - \mathbf{x}') C(\mathbf{x}') + \frac{\gamma}{2} \int d^2x |C(\mathbf{x}) - C(\hat{\mathbf{R}} \cdot \mathbf{x})|^2 \right], \quad (399)$$

ture ($\gamma = 0$). Given images of a certain size, and a known kernel K , there is a limit to the smallest value of γ that can be distinguished reliably from zero, and we can compare this statistical limit to the performance of human observers. This is more or less what Barlow did, although he used blurred random dots rather than the Gaussian


textures considered here; the idea is the same, and all the details become the same in the limit of many dots. The result is that human observers come within a factor of two of the statistical limit for detecting γ or its analog in the random dot patterns. [Show a 1D version of this problem in a figure.]

One can use similar sorts of visual stimuli to think about motion, where rather than having to recognize a match between two halves of a possibly symmetric image we have to match successive frames of a movie. Here again human observers can approach the statistical limits, as long as we stay in the right regime: we seem not to make use of fine dot positioning (as would be generated if the kernel K only contained low order derivatives) nor can we integrate efficiently over many frames. These results are interesting because they show the potentialities and limitations of optimal visual computation, but also because the discrimination of motion in random movies is one of the places where people have tried to make close links between perception and neural activity in the (monkey) cortex.

Let us look in detail at the case of visual motion estimation, using not humans or monkeys, but a smaller system which we have met once before—the visual system of the fly, which we have met already in Section I.A. If you watch a fly flying around in a room or outdoors, you will notice that flight paths tend to consist of rather straight segments interrupted by sharp turns and acrobatic interludes. These observations can be quantified through the measurement of trajectories during free flight, and in experiments where the fly is suspended from a torsion balance or a fine tether. Given the aerodynamics for an object of the fly's dimensions, even flying straight is tricky. In the torsion balance one can demonstrate directly that motion across the visual field drives the generation of torque, and the sign is such as to stabilize flight against rigid body rotation of the fly. Indeed one can close the feedback loop by measuring the torque which the fly produces and using this torque to (counter)rotate the visual stimulus, creating an imperfect 'flight simulator' for the fly in which the only cues to guide the flight are visual; under natural conditions the fly's mechanical sensors play a crucial role. Despite the imperfections of the flight simulator, the tethered fly will fixate small objects, thereby stabilizing the appearance of straight flight. Similarly, aspects of flight behavior under free flight conditions can be understood if flies generate torques in response to motion across the visual field, and that this response is remarkably fast, with a latency of just ~ 30 msec. The combination of free flight and torsion balance experiments strongly suggests that flies can estimate their angular velocity from visual input alone, and then produce motor outputs based on this estimate.

Voltage signals from the receptor cells are processed by several layers of the brain, each layer having cells organized on a lattice which parallels the lattice of lenses

visible from the outside of the fly. As shown in Fig 59, after passing through the lamina, the medulla, and the lobula, signals arrive at the lobula plate. Here there is a stack of about 50 cells which are sensitive to different components of motion. These cells have imaginative names, such as H1 and V1, which respond to horizontal and vertical components of motion, respectively. If one kills individual cells in the lobula plate then the simple experiment of moving a stimulus and recording the flight torque no longer works, strongly suggesting that these cells are an obligatory link in the pathway from the retina to the flight motor. Taken together, these observations support a picture in which the fly's brain uses photoreceptor signals to estimate angular velocity, and encodes this estimate in the activity of a few neurons.⁴⁸ What



de Ruyter van Steveninck and Bialek, FIG 3.

FIG. 59 The visual system of a fly, from the retina to the motion sensitive cells of the lobula plate. From de Ruyter van Steveninck & Bialek (2002).

⁴⁸ You should be skeptical of any claim about what the brain computes, or more generally what problems an organism has to solve in order to explain some observed behavior. The fact that flies can stabilize their flight using visual cues, for example, does *not* mean that they compute motion in any precise sense—they could use a form of 'bang-bang' control that needs knowledge only of the algebraic sign of the velocity, although I think that the tor-




FIG. 60 The limits to motion detection. At top, a possible pattern of contrast (normalized light intensity) vs. position or angle in the visual world. Blue denotes the original pattern, and green illustrates a shift by one tenth of the spacing between photoreceptors. The second panel from the top shows the blurring and sampling of the image, with Gaussian apertures that provide a model for the optics of the fly’s eye. Note that the spacing between photoreceptors is comparable to width of the diffraction blur. The third panel shows the signal arriving at each photoreceptor. We see that the blurring reduces the contrast enormously. The bottom panel illustrates the effect of adding noise, here with an amplitude expected if each snapshot involves counting an average of 10^3 photons. Insets show the distribution of signals plus noise in response to the original (blue) and shifted (green) images. Despite the large differences between the two initial patterns, only one of the five receptor cells shown here would be able to come near to reliable detection. The experiments described in the text are done under conditions of even smaller signal-to-noise ratios.

can we say about the physical limits to the precision of this computation?

Suppose that we look at a pattern of typical contrast C and it moves by an angle $\delta\theta$, as schematized in Fig 60. A single photodetector element will see a change in contrast of roughly $\delta C \sim C \cdot (\delta\theta/\phi_0)$, where ϕ_0 is the angular scale of blurring due to diffraction. If we can measure for a time τ , we will count an average number of photons $R\tau$, with R the counting rate per detector, and hence the noise can be expressed as a fractional precision in intensity of $\sim 1/\sqrt{R\tau}$. But fractional intensity is what we mean by contrast, so $1/\sqrt{R\tau}$ is really the contrast

noise in one photodetector. To get the signal-to-noise ratio we should compare the signal and noise in each of the N_{cells} detectors, then add the squares if we assume (as for photon shot noise) that noise is independent in each detector while the signal is coherent:

$$SNR \sim N_{\text{cells}} \cdot \left(\frac{\delta\theta}{\phi_0}\right)^2 C^2 R\tau. \quad (400)$$

Motion discrimination is hard for flies because they have small lenses and hence blurry images (ϕ_0 is large) and because they have to respond quickly (τ is small); typical photon counting rates in a laboratory experiment are $R \sim 10^4 \text{ s}^{-1}$ and outside on a bright day one can get to $R \sim 10^6 \text{ s}^{-1}$. Under reasonable laboratory conditions—and taking account of all the factors that go in front of our rough Eq (400) in a more careful calculation—the optimal estimator would reach $SNR = 1$ at an angular displacement of $\delta\theta \sim 0.05^\circ$.

We can test the precision of motion estimation in two very different ways. One is similar to the experiments we have discussed already, where we are forced to choose between two alternatives and measure the reliability of this choice. A single neuron responds to sudden steps of motion with a brief volley of action potentials which we




FIG. 61 Motion discrimination with the fly’s H1 neuron, from de Ruyter van Steveninck & Bialek (1995). At left, a schematic of the spikes in response to a transient stimulus, such as a step of motion. We can describe the response by the time until the first spike τ_0 , the time from the first spike to the second τ_1 , Alternatively we can just count the spikes that have occurred up to a certain time after the stimulus, or we could at some fixed time resolution describe the whole pattern of spikes as a binary word. In each case we can analyze the discriminability of different stimuli by accumulating, over many repeated presentations of each stimulus, the distribution of responses. At right, an example of this analysis, focusing on the single interspike interval τ_1 in response to steps that differ in size by 0.12° . Long intervals correspond to the weaker stimulus, and from the cumulative probability distributions in the top panel we can read off the probabilities of correct identification of each stimulus.

sion balance experiments argue against such a model. It also is a bit mysterious why we find neurons with such understandable properties: one could imagine connecting photoreceptors to flight muscles via a network of neurons in which there is nothing that we could recognize as a motion-sensitive cell. Thus it is not obvious either that the fly must compute motion or that there must be motion-sensitive neurons.

can label as occurring at times t_1, t_2, \dots . We as observers of the neuron can look at these times and try to decide whether the motion had amplitude θ_+ or θ_- ; the idea is exactly the same as in earlier discussions of discrimination of signal vs noise, but here we have to measure the relevant probability distributions rather than making assumptions about their form; see Fig 61. Doing the integrals, one finds that looking at spikes generated in the first ~ 30 msec after the step (as in the fly's behavior) we can reach the reliability expected for $SNR = 1$ at a displacement $\delta\theta = |\theta_+ - \theta_-| \sim 0.12^\circ$, within a factor of two of the theoretical limit set by noise in the photodetectors.

It is worth noting a few more points that emerge from Fig 61 and further analyses of this experiment. First, on the ~ 30 msec time scale of relevance to behavior, there are only a handful of spikes. This is partly what makes it possible to do the analysis so completely, but it also is a lesson for how we think about the neural representation of information in general. Second, we can dissect the contributions of individual spikes to show that each successive spike makes a nearly independent contribution to the signal to noise ratio for discrimination, so there is essentially no redundancy. Finally, the motions we are discussing—motions close to the physical limits of detectability, and motions that real neurons can represent reliably—are much smaller than the lattice spacing on the retina or the nominal “diffraction limit” of angular resolution $\sim 1^\circ$. Analogous phenomena have been known in human vision for more than a century, and are called hyperacuity.

The step discrimination experiment gives us a very clear view of reliability in the neural response, but as with the other discrimination experiments discussed above it's not a very natural task. An alternative is to ask what happens when the motion signal (angular velocity $\dot{\theta}(t)$) is a complex function of time. Then we can think of the signal to noise ratio in Eq. (400) as being equivalent to a spectral density of displacement noise $N_\theta^{\text{eff}} \sim \phi_0^2 / (N_{\text{cells}} C^2 R)$, or a generalization in which the photon counting rate is replaced by an effective, frequency dependent, rate related to the noise characteristics of the photoreceptors, as in Fig 13. It seems likely, as discussed above, that the fly's visual system really does make a continuous or running estimate of the angular velocity, and that this estimate is encoded in the sequence of discrete spikes produced by neurons like H1. It is not clear that any piece of the brain ever “decodes” this signal in an explicit way, but if *we* could do such a decoding we could test directly whether the accuracy of our decoding reaches the limiting noise level set by noise in the photodetectors.

Decoding spike trains, at least under certain conditions, is much easier than one might have expected. The idea, shown in Fig 62, is that each spike contributes a small transient blip to our estimate of the signal vs. time, and to obtain the full estimate we add up all these small




FIG. 62 Decoding continuous motion signals from spikes generated by the H1 neuron, from Bialek et al (1991). At left, dashed curve indicates the true stimulus, angular velocity as a function of time; solid line is the result of the decoding process, from Eq (403). Tick marks below the stimulus indicate the spikes generated in a single presentation of this stimulus (downward ticks) or its negative (upward ticks). This consideration of a hypothetical neuron that sees the negative stimulus is meant to restore symmetry between positive and negative velocities and corresponds roughly to the response of the H1 neuron on the other side of the fly's head, which has the opposite direction selectivity. At right is the spectral density of errors in the reconstruction. The error is reported as a displacement error, so the spectrum grows as $1/\omega^2$ for low frequencies. Also shown is the spectrum of the stimulus (smooth line) and the limiting noise level computed from the actual noise levels measured in fly photoreceptors under the same conditions as for these experiments on H1. Reconstruction error and the physical limit to precision converge at high frequencies, so that the fly approaches optimal performance.

contributions. Thus, if the signal we are interested in is $s(t)$, our estimate is

$$s_{\text{est}}(t) = \sum_i f(t - t_i), \quad (401)$$

where t_i are the spike arrival times as before, and we can choose the filter $f(t)$ to minimize the errors

$$\chi^2 \equiv \int dt \left| s(t) - s_{\text{est}}(t) \right|^2. \quad (402)$$

Like most neurons, H1 has a sign preference for its inputs—motion in one direction generate more spikes, while motion in the opposite direction generates fewer spikes. Thus, large negative velocities cause H1 to go silent, and in these periods we would have no basis for inferring the detailed waveform of velocity vs. time. Fortunately, the fly has two H1 neurons, one on each side of

the head, with opposite direction preferences. We could record from both cells, or we could use the fact that the two cells see opposite motions relative to their own preference, and look at the responses of one neuron to both a stimulus and the opposite motion. If the spikes in these two cases are $\{t_i^+\}$ and $\{t_i^-\}$, we can make a more symmetric reconstruction

$$s_{\text{est}}(t) = \sum_i [f(t - t_i^+) - f(t - t_i^-)]. \quad (403)$$

Again, we choose the filter $f(t)$ to minimize χ^2 .⁴⁹

In Figure 62 we see that the reconstruction of the velocity waveform in fact is quite accurate. More quantitatively, the power spectrum of the errors in the reconstructed signal approaches the limit set by noise in the photoreceptor cells, within a factor of two at high frequencies. Further, one can change, for example, the image contrast and show that the resulting error spectrum scales as expected from the theoretical limit.

To the extent that the fly's brain can estimate motion with a precision close to the theoretical limit, we know that the act of processing itself does not add too much noise. But being quiet is not enough: to make maximally reliable estimates of nontrivial stimulus features like motion one must be sure to do the correct computation. Making this idea precise is in the same spirit as the discussion, in Section I.D, of pooling single photon signals from multiple rod cells at the level of bipolar cells. There we saw how the different orders of nonlinearity and summation result in very different final signal-to-noise ratios, even though all we are trying to do is add. Here the problem is more difficult, because the fly wants to estimate a feature of the visual world which is not directly reflected in the signals of any single receptor cell.

Problem 71: (Relatively) simple estimation problems.

Suppose that someone draws a random number x from a probability distribution $P(x)$. Rather than seeing x itself, you get to see only a noisy version, $y = x + \eta$, where η is drawn from a Gaussian distribution with variance σ^2 , so that

$$P(y|x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2\sigma^2}(y-x)^2\right]. \quad (404)$$

Having seen y , your job is to estimate x .

(a.) Show that everything you know about x by virtue of observing y can be written in a way that suggests an analogy with statistical mechanics,

$$P(x|y) = \frac{1}{Z(y)} \exp\left[-\frac{V_{\text{eff}}(x)}{k_B T_{\text{eff}}} + \frac{F_{\text{eff}} x}{k_B T_{\text{eff}}}\right], \quad (405)$$

where

$$\frac{V_{\text{eff}}(x)}{k_B T_{\text{eff}}} = -\ln P(x) + \frac{x^2}{2\sigma^2} \quad (406)$$

$$k_B T_{\text{eff}} = \sigma^2 \quad (407)$$

$$F_{\text{eff}} = y. \quad (408)$$

(b.) From the discussion in Section I.D, we know that if we define “best” to be the estimator that minimizes χ^2 , then the best estimator is the conditional mean,

$$x_{\text{est}}(y) = \int dx x P(x|y). \quad (409)$$

Construct $x_{\text{est}}(y)$ in the case where $P(x)$ is a Gaussian with unit variance. Show that this estimate, although “best,” is systematically wrong. That is, if we average $x_{\text{est}}(y)$ over the distribution $P(y|x)$, we do not recover x itself. Explain why this can still be the best estimate.

(c.) Now consider the case $P(x) = (1/2)\exp(-|x|)$. Show that, even though the transformation from what we are interested in (x) to what we measure (y) is linear, the optimal estimator is nonlinear. In particular, if rather than asking for an estimator that minimizes χ^2 , we ask for the most probable value of x given y , show that the optimal estimator involves a threshold nonlinearity.

Motion estimation is an example of the more general problem of perceptual estimation. The data to which the brain has access are the responses of receptor cells, and the goal is to estimate some feature of the world. The first key step is to use Bayes' rule, combining the noisy data from the receptors with our prior knowledge that some things are more likely than others. Schematically,

$$\begin{aligned} &P(\text{feature}|\text{receptor responses}) \\ &= \frac{P(\text{receptor responses}|\text{feature})P(\text{feature})}{P(\text{receptor responses})}. \end{aligned} \quad (410)$$

The second key step is to note that receptors typically don't respond directly to the features of interest, but rather to raw sensory signals such as light intensity, sound pressure in the auditory system, the concentrations of specific molecular species in complex odors, etc.. Continuing schematically, let's denote the full spatiotemporal pattern of light intensities falling on the retina by \mathcal{I} . Receptor responses really depend on \mathcal{I} , which in turn is correlated with the feature that we want to estimate. Thus,

$$P(\text{receptor responses}|\text{feature}) = \int D\mathcal{I} P(\text{receptor responses}|\mathcal{I})P(\mathcal{I}|\text{feature}), \quad (411)$$

and putting all the terms together we have

$$P(\text{feature}|\text{receptor responses}) = \frac{1}{P(\text{receptor responses})} \int DI P(\text{receptor responses}|\mathcal{I})P(\mathcal{I}, \text{feature}). \quad (412)$$

If the lights are bright, and the noise level in the photoreceptors is low, it is plausible that knowing the pattern of receptor responses is almost equivalent to knowing the spatiotemporal pattern of light intensities \mathcal{I} , and hence viewed as a function of \mathcal{I} the distribution $P(\text{receptor responses}|\mathcal{I})$ is very sharply peaked. Then the entire structure of the optimal computation that maps receptor responses to the desired feature is controlled by $P(\mathcal{I}, \text{feature})$, which is a property of the world that we live in rather than of our eyes or brains. This is perhaps our most important qualitative conclusion: optimal estimates of sensory features involve computations determined by the structure of the world around us. To the extent that our brains, and those of other animals, make optimal estimates, this means that the way in which we process the world is set by the physics of our environment, not by peculiarities of our biological hardware.

For the case of motion estimation, what is the structure of $P(\mathcal{I}, \text{feature})$? For simplicity let's think about a one-dimensional version of the problem, so that spatiotemporal pattern of light intensity $\mathcal{I} \equiv I(x, t)$. Then if a small piece of the visual world is moving rigidly relative to us with a velocity v , we should have $I(x, t) = I_0(x - vt)$. Then we can take derivatives in space and time,

$$\frac{\partial I(x, t)}{\partial x} = I_0'(x - vt) \quad (413)$$

$$\frac{\partial I(x, t)}{\partial t} = -vI_0'(x - vt). \quad (414)$$

Thus, we can compute the velocity as a ratio of spatial and temporal derivatives,

$$v_{\text{est}} = -\frac{\partial I(x, t)/\partial t}{\partial I(x, t)/\partial x}. \quad (415)$$

This is correct, but we have derived it by pushing to extremes. First we said that noise in the receptor responses is negligible, so we can say that we are effectively computing functions of the light intensity itself. Then we assumed that the dynamics of the light intensity is determined only by motion at the single velocity v . If either of these assumptions breaks down, our “gradient based” estimator of velocity, Eq (415) gets into serious trouble.

When we deal with noisy data we develop several intuitions. First, the nature of our measurements is such that there usually is relatively more noise at higher frequencies, both in time and in space. Thus, to suppress noise, we average. Conversely, if we differentiate, we expect that noise will be amplified, since differentiation enhances higher frequencies. Second, when we have a noisy

measurement, it is dangerous to put this in the denominator of a ratio—there is a chance that we will divide by zero, because of a fluctuation. The gradient based estimator compounds these sins, differentiating and then taking a ratio. We expect that this will be a disaster if our low noise assumptions are violated.

Problem 72: Ratios of noisy numbers. Suppose that we have two numbers that we try to measure, a and b . Our measurements, which we can call \hat{a} and \hat{b} , give us the values of a and b but with some added Gaussian noise, so that

$$P(\hat{a}|a) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(\hat{a}-a)^2/2\sigma^2}; \quad (416)$$

for simplicity we'll assume that the noise level is the same for our measurements of b , so that

$$P(\hat{b}|b) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(\hat{b}-b)^2/2\sigma^2}. \quad (417)$$

What we would like to do is to estimate the ratio $r \equiv a/b$ from our measurements \hat{a} and \hat{b} .

(a.) Suppose we make form a naive estimate just by taking the ratio of our measurements, $r_{\text{est}}^{\text{naive}} = \hat{a}/\hat{b}$. Do a small simulation to examine numerically the probability distribution of this estimate. In particular, consider the case where $a = b = 1$, so the correct answer is $r = 1$. If $\sigma = 0.1$, presumably $r_{\text{est}}^{\text{naive}}$ stays close to this correct answer, but what happens at $\sigma = 0.2$ or 0.5 ? How does the variance of the estimator $r_{\text{est}}^{\text{naive}}$ change as the noise level σ increases? Be sure to check in your simulation that you have enough samples to get a reliable measure of the variance. Is there anything suspicious in this computation, especially at larger σ ?

(b.) Look more closely at the right hand tail of the distribution of $r_{\text{est}}^{\text{naive}}$, that is the behavior of $P(r_{\text{est}}^{\text{naive}} \gg 1)$ in the case where $a = b = 1$. Plot your numerical results on linear, semilog, and log-log plots to see if you can recognize the shape of the tail. If the shape changing with the noise level σ ? Try to make a precise statement based on your simulations. I have left this somewhat open ended.

(c.) Try to derive analytically the regularities that you found in [b].

(d.) Although we think of \hat{a} and \hat{b} as measurements of the separate variables a and b , really all we want to know is the ratio $r \equiv a/b$. Show that the best estimate can be written, using Bayes' rule, as

$$r_{\text{est}}(\hat{a}, \hat{b}) = \int dr \frac{r}{P(\hat{a}, \hat{b})} \int da \int db \delta\left(r - \frac{a}{b}\right) P(\hat{a}|a)P(\hat{b}|b)P(a, b). \quad (418)$$

Make as much progress as you can evaluating these integrals on the hypothesis that the prior distribution $P(a, b)$ is broad and featureless. If you want to proceed analytically, you may find it useful to introduce a Fourier representation of the delta function, and look for a saddle point approximation. Numerically, you could assume, for example, that $P(a, b)$ is uniform over some region of the $a - b$ plane, and just do the integrals for representative values of \hat{a} and \hat{b} , mapping the function $r_{\text{est}}(\hat{a}, \hat{b})$. Can you verify that $r_{\text{est}}^{\text{naive}}$ is close to optimal at very small values of σ ? What happens at larger values of σ ? If σ is fixed, what happens as $b \rightarrow 0$?

The most obvious problem with the gradient motion estimator in Eq (415) is simply that it is not well defined when the spatial derivative becomes small. This problem exists even if noise in the photoreceptors is small. To address the problem we have to understand what the distribution $P(\mathcal{I}, \text{feature})$ looks like. Conceptually, what we want to do is simple. Imagine taking a walk on a very still day, so that motions of the world relative to our retina (or relative to the fly's retina) are dominated by our own motion. If we carry a camera as we walk, we can take a movie, and we can also put a gyroscope on the camera to monitor its motion. What emerges from such an experiment, then, is a set of samples drawn out of the distribution $P(\mathcal{I}, \text{feature})$. In particular, pixel by pixel and moment by moment, we can compute the spatial and temporal derivatives in the movie, and measure the velocity as well, so that we sample the distribution $P(\partial I/\partial t, \partial I/\partial x, v)$.

If the gradient based estimate of motion were exact, then the distribution $P(\partial I/\partial t, \partial I/\partial x, v)$ would be very sharply peaked along a ridge where $v = -(\partial I/\partial t)/(\partial I/\partial x)$. To see if this is right, we can compute directly the optimal estimator. We know that the best estimate in the sense of χ^2 is the conditional mean, so should compute⁵⁰

$$v_{\text{est}}(\partial_t I, \partial_x I) = \int dv v \frac{P(\partial_t I, \partial_x I, v)}{P(v)}. \quad (419)$$

The results of this computation, based on a walk in the woods, are shown in Fig 63.⁵¹ We see that, when the spatial gradients are large, the contours of constant v_{est} really are straight lines, as expected from the gradient based estimator. But when the spatial gradients are smaller, a new structure emerges, which is more closely approximated by a *product* of derivatives, $v_{\text{est}} \propto (\partial I/\partial t) \times (\partial I/\partial x)$, rather than a ratio. As you can see in the following problem, the same product structure emerges if we go back to the general formulation and take the limit of high noise levels.

Problem 73: Series expansion of the optimal estimator at low signal-to-noise ratios. We know from Section I.A that

⁵⁰ I need to make a segue between the notation $\partial I/\partial x$ and $\partial_x I$.

⁵¹ Although conceptually simple, to generate Fig 63 requires measuring light intensities with spatial and temporal resolution matched to that of the retina, but collecting much more light so that photon shot noise in these measurements will be less than that in the retina and one can meaningfully claim to measure intensity at the input to the visual system. For details, as always, see the references at the end of the section.




FIG. 63 Optimal estimates of angular velocity as a function of local spatial and temporal gradients of light intensity. Computed from the theory described in the text, with the joint distribution of movies and motion sampled experimentally. Images are collected through an optical system that matches the fly's eye, and smoothed in time with a filter that optimizes estimation performance. At small signals, near the center of the plot, we see that moving along a line of constant physical velocity (in white; $\partial_t I + v \partial_x I = 0$) results in a changing estimate—a systematic error; only for large signals is the optimal estimate veridical. Experiments by SR Sinha & RR de Ruyter van Steveninck.

photoreceptors in the fly respond linearly to changes in light intensity or contrast [point back to specific equations; check consistency of notation]. If the fly is rotating relative to the world along an angular trajectory $\theta(t)$, then the spatiotemporal pattern of contrast (again in a one-dimensional model) is $C(x - \theta(t), t)$. Individual cells respond with voltages $V_n(t)$ given by

$$V_n(t) = \int dt' T(t - t') \int dx M(x - x_n) C(x - \theta(t'), t'), \quad (420)$$

where $T(\tau)$ is the temporal impulse response function and $M(x - x_n)$ is an aperture function centered on a lattice point x_n in the retina.

(a.) Show that the distribution of all the voltages given the trajectory can be written as

$$P[\{V_n(t)\}|\theta(t)] \propto \int DC P[C] \exp \left[-\frac{1}{2} \sum_n \int \frac{d\omega}{2\pi} \frac{|\tilde{V}_n(\omega) - \langle \tilde{V}_n(\omega) \rangle|^2}{N_V(\omega)} \right] \quad (421)$$

where the mean voltages are, in the Fourier representation,

$$\langle \tilde{V}_n(\omega) \rangle = \tilde{T}(\omega) \int dx M(x - x_n) \int dt e^{+i\omega t} C(x - \theta(t), t), \quad (422)$$

$N_V(\omega)$ is the power spectrum of the voltage noise, and $P[C]$ is the distribution of contrast that the fly would observed if held at $\theta = 0$.

(b.) The optimal estimator is the conditional mean,

$$\hat{\theta}_{\text{est}}(t_0) = \int D\theta \hat{\theta}(t_0) P[\theta(t)|\{V_n(t)\}] \quad (423)$$

$$P[\theta(t)|\{V_n(t)\}] = \frac{P[\{V_n(t)\}|\theta(t)]P[\theta(t)]}{P[\{V_n(t)\}]} \quad (424)$$

Evaluate all the integrals in a perturbation series, assuming that the average voltage responses are small compared with the noise

level. You should find that the leading term is

$$\dot{\theta}_{\text{est}}(t) \approx \sum_{nm} \int d\tau \int d\tau' V_n(t - \tau) K_{nm}(\tau, \tau') V_m(t - \tau'). \quad (425)$$

Relate the kernel $K_{nm}(\tau, \tau')$ to expectation values in the distributions $P[C(x, t)]$ and $P[\theta(t)]$.

(c.) Can you reformulate the expansion so that instead of expanding for small overall signal-to-noise ratio (small R), you expand for small instantaneous signals, that is for small $V_n(t)$? What happens to the kernels in this case? It seems obvious that there shouldn't be a linear term in this expansion. Can there be a third order term? If such a term exists, what happens to the optimal estimate of velocity when if we show the same movie, but with inverted contrast (exchanging black for white)?

We can understand the low signal to noise ratio limit by realizing that when something moves there are correlations between what we see at the two space-time points (x, t) and $(x + v\tau, t + \tau)$. These correlations extend to very high orders, but as the background noise level increases the higher order correlations are corrupted first, until finally the only reliable thing left is the two-point function, and closer examination shows that near neighbor correlations are the most significant: we can be sure something is moving because signals in neighboring photodetectors are correlated with a slight delay. This form of “correlation based” motion computation, schematized in Fig 64, was suggested long ago by Reichardt and Hasenstein based on behavioral experiments with beetles.

There are two clear signatures of the correlation model. First, since the receptor voltage is linear in response to image contrast, the correlation model confounds contrast with velocity: all things being equal, doubling the image contrast causes our estimate of the velocity to increase by a factor of four (!). This is an observed property of the flight torque that flies generate in response to visual motion, at least at low contrasts, and the same quadratic behavior can be seen in the rate at which motion sensitive neurons generate spikes, as shown in Fig 65. Even humans experience the illusion of contrast dependent motion perception at very low contrast. Although this might seem strange, it's been known for decades.

The second signature of correlation computation is that we can produce movies which have the right spatiotemporal correlations to generate a nonzero estimate $\dot{\theta}_{\text{est}}$ but don't really have anything in them that we would describe as “moving” objects or features. Consider a spatiotemporal white noise movie $\psi(\mathbf{x}, t)$,

$$\langle \psi(\mathbf{x}, t) \psi(\mathbf{x}', t') \rangle = \delta(\mathbf{x} - \mathbf{x}') \delta(t - t'), \quad (426)$$

and then add the movie to itself with a weight and an offset:

$$C(\mathbf{x}, t) = \psi(\mathbf{x}, t) + a\psi(\mathbf{x} + \Delta\mathbf{x}, t + \Delta t). \quad (427)$$

Composed of pure noise, there is nothing really moving here. If you watch the movie, however, there is no question that you think it's moving, and the fly's neurons respond too (just like yours, presumably). Even more impressive is that if you change the *sign* of the weight a , then the direction of motion reverses, as predicted from the correlation model.

Problem 74: Motion from correlations alone. Generate the image sequences described in the previous paragraph, and verify that you (and your friends) perceive them as moving.

(a.) Play with the amplitude and sign of the weight a to see how it influences your perception. Can you find a regime in which the speed of motion seems to depend on $|a|$? Can you verify the reversal of motion when $a \rightarrow -a$?

(b.) Compute the correlation function $\langle C(\mathbf{x}, t) C(\mathbf{x}', t') \rangle$; for simplicity you might want to confine your attention to a one dimensional example. Consider also the correlation function for a genuine moving image, in which $C(\mathbf{x}, t) = C_0(\mathbf{x} - \mathbf{v}t)$. If $\mathbf{v} = \Delta\mathbf{x}/\Delta t$, how do the two correlation functions compare?




FIG. 64 The correlator model of visual motion detection, adapted from Reichardt (1961). A spatiotemporal contrast pattern $C(x, t)$ is blurred by the photoreceptor point spread function, $M(x)$, and sampled by an array of photoreceptors, two of which (neighboring photoreceptors numbers $n - 1$ and n) are shown here. After phototransduction, the signals in each photoreceptor are filtered by two different linear filters, $f(t)$ and $g(t)$. The outputs of these filters from the different photoreceptors, $s_1(t)$ and $s_3(t)$ from photoreceptor n and $s_2(t)$ and $s_4(t)$ from photoreceptor $n - 1$ are multiplied and one of these products is subtracted from the other by the addition unit, yielding a direction selective response. Thanks to Rob de Ruyter for this figure.




FIG. 65 Responses of the H1 neuron to moving scenes with varying contrast. Scenes consist of bars with random intensities, moving at constant velocity. At left, at one particular velocity we measure the rate at which H1 generates action potentials, as a function of contrast. Lower panel expands the region at low contrast, emphasizing the quadratic behavior. At right, the responses at multiple velocities, showing that the “saturated” response at high contrast still is sensitive to the speed of movement. [does this appear in a paper? details of stimuli?] Thanks to Rob de Ruyter for this figure.

The optimal motion estimator illustrates the general tradeoff between systematic and random errors. If we really are viewing an image that moves rigidly, so that $C(x, t) = C(x + vt)$, then there is no question that the “right answer” is to compute v as the ratio of temporal and spatial derivatives. Any departure from this involves making a systematic error. But, as discussed above, taking derivatives and ratios are both operations which are perilous in the presence of noise. To insulate the estimate from random errors driven by such noise (or, more generally, by aspects of the image dynamics that are not related to motion), we must calculate something which, typically, will not give the “right answer” even on average—we accept some systematic errors in order to reduce the impact of random errors. In the context of perception, systematic errors have a special name: illusions.

Could the theory of optimal estimation be a quantitative theory of illusions, grounded in physical principles? Colloquially, we say that “to err is human,” and it is conventional to assume that cases in which biological systems get the wrong answer to their signal processing problems provide evidence regarding the inadequacies of the biological hardware. Is it possible that, rather than being uniquely human or biological, to err is the optimal response to the limits imposed by the physical world?

The long history of the correlation model provides ample testimony that insect visual systems make the kind of systematic errors expected from the optimal estimator, but precisely because of this long history it is hard to view these as successful predictions. It would be more compelling if we could show that the same system which is well described by the correlator under some conditions crosses over to something more like the ratio of derivatives model at high signal-to-noise ratio, but this has been elusive. The contrast dependence of the response in the motion sensitive neurons saturates at high contrast, and this saturated response still varies with velocity (Fig 65), as if the larger signals allow the system to disentangle ambiguities and recover a veridical estimate, but other experiments suggest that errors inherent in the correlation model persist even with strong signals. Humans easily see the illusion of motion with the noise movies of Eq (427), as well as other motion illusions, but at high signal-to-noise ratios our visual systems recover estimates of velocity which are not systematically distorted, suggesting that in primates there is some sort of crossover between different limits of the motion computation, and there are efforts to make the correspondence with the optimal estimator more quantitative. Experiments under more natural, free flight conditions show that both flies and bees have access to veridical estimates of their translational velocity and can use this to control their flight speed, in contrast to what one would have expected from the correlator model, and it worth noting that the responses of the motion-sensitive neurons are also very different under more natural conditions.

[This needs to be clearer] In Figure 66 we see the responses of the H1 neuron to the rotation of a fly, outside under nearly natural conditions. During the course




FIG. 66 Responses of the H1 neuron as a fly is rotated outside, over a period during which the mean light level is falling. [fill in the caption. does this appear in a paper? details of stimuli?] Thanks to Rob de Ruyter for this figure.

of the experiment, the sun was going down, and so the mean light level varied by several orders of magnitude as the same trajectory of angular velocity vs. time was repeated over and over. The integral of the trajectory was not quite zero, however, so that on each repetition the spatial pattern of light intensity was a bit different even if the angular velocity was the same. At the start of the experiment, the responses are extremely vigorous, and insensitive to the variations in the spatial structure of the visual environment. As the light level falls, responses become weaker, but more dramatically we see that there is a systematic variation from repetition to repetition, which appears as a diagonal pattern of spikes across the upper part of Fig 66. Thus, when signal-to-noise ratios are high in the natural environment, H1 responds to time dependent velocities and largely ignores the spatial structure of its environment, while at lower signal-to-noise ratios the confounding of spatial structure and motion becomes more and more obvious. This pattern is in agreement with the expectations from optimal estimation theory, according to which such systematic errors arise only from the need to insulate the computation from random noise.

What we would really like is to have methods of dissecting the computation that has been done by a neuron, simply by analyzing the relationship between visual inputs and spiking outputs under natural conditions. This is a huge challenge, and obviously would be interesting in many other contexts. Approaches to this problem are discussed in Appendix A.7, where we also see results that come closest to a smoking gun for the crossover between correlator and gradient computations.

For visual signal processing, getting our hands on the true distribution of signals in the natural environment is a difficult experiment. For seemingly more complex “cognitive” judgments, the situation, perhaps surprisingly, is much simpler. To give an example, suppose that you are told of a member of the United States Congress who has served for $t = 15$ years. What is your prediction for how long his total term will last? To keep things as simple as possible, let’s assume you are not told anything about the politics of this congressman or his district; all you have to work with is $t = 15$ and your general knowledge of the turnover of elected officials. Obviously your knowledge is probabilistic, so we use Bayes’ rule to write

$$P(t_{\text{total}}|t) \propto P(t|t_{\text{total}})P(t_{\text{total}}). \quad (428)$$

If the moment at which the question is asked is not somehow synchronized to the length of congressional terms, then we have to assume that $P(t|t_{\text{total}})$ is uniform, $P(t|t_{\text{total}}) = 1/t_{\text{total}}$. Thus our inference is controlled by the “prior” distribution $P(t_{\text{total}})$, and we can look this up in a database about the history of the congress. Finally, if you must pick one value of t_{total} , it makes sense in this context to choose the median, the point at which the actual value of t_{total} is equally likely to be longer or shorter

than your estimate. As an example, if $P(t_{\text{total}})$ is a reasonably narrow Gaussian distribution, then for t much less than the mean $\langle t_{\text{total}} \rangle$, our best estimate of t_{total} is just $\langle t_{\text{total}} \rangle$ itself, while if the time t is much larger than the mean then our best estimate is only slightly higher than t , which makes sense. Other priors, of course, can give qualitatively different results.

Problem 75: Estimating t_{total} . Derive the results just stated for the Gaussian prior. Consider also cases where $P(t_{\text{total}}) \propto t_{\text{total}}^{-\gamma}$ or $P(t_{\text{total}}) \propto t_{\text{total}}^n e^{-t_{\text{total}}/\tau}$.

The example of congressional terms is not unique. We could ask, as insurance companies do (albeit with more input data), about human lifespans: if you meet someone of age t , what is your best guess about their life expectancy? If you make a phone call and have been on hold for t minutes, what is your best guess about the total time you will have to wait? If you find yourself on line t of a poem, what is your best guess about the total length of text? Nor is the structure of the problem bound to time, as such: suppose you learn that a movie has collected t dollars in gross receipts; what is your best guess about what its total earnings will be? All these problems have in common that we can look up the correct distribution $P(t_{\text{total}})$. Another important feature is that we can just go ask people what they think, and see how they do relative to the predictions for optimal estimation based




FIG. 67 Estimation of totals based on one observation, from Griffiths & Tenenbaum (2006). The top row shows the priors $P(t_{\text{total}})$ measured from real world data. The bottom panel compares people’s predictions (points) based on one observation t with the optimal median estimator (solid lines) and a naive “universal” estimate $\hat{t}_{\text{total}} = 2t$. For the reigns of Pharaohs and the telephone waiting times, dashed lines show optimal estimators for $P(t_{\text{total}}) \propto t_{\text{total}} e^{-t_{\text{total}}/\tau}$ ($\tau = 17.9$) and $P(t_{\text{total}}) \propto t_{\text{total}}^{-\gamma}$ ($\gamma = 2.43$), respectively.

on the priors appropriate to our real world. The results from such an experiment are shown in Fig 67.

I found the results of Fig 67 quite astonishing when I first saw them. The time it takes to bake a cake comes from a very irregular distribution, but people seem to know this distribution and estimate accordingly. They are a bit confused about how long the Pharaohs reigned, but their confusion is consistent: estimation of t_{total} behaves as if the subjects know the shape of $P(t_{\text{total}})$ but are off on the mean time scale, and if you ask another group of subjects to guess the mean reign of the Pharaohs, they deviate from the right answer by the same factor. Important as the telephone problem may be, this is one case where there is no convenient data to which we can refer, so this case remains untested. In all the other cases, however, spanning seemingly very different domains of knowledge and very different shapes for $P(t_{\text{total}})$, people are performing close to the optimum.

If we trace through the details of optimal estimation theory, one can see that construction of the correct estimator involves knowing not only the distribution of signals, but also the distribution of noise. Perhaps the simplest illustration of this is given by the problem of combining two measurements. Suppose that we are interested in x , but we observe

$$y_1 = x + \eta_1 \quad (429)$$

$$y_2 = x + \eta_2, \quad (430)$$

where the noise levels on the two measurements are generally different, $\langle \eta_1^2 \rangle = \sigma_1^2$ and $\langle \eta_2^2 \rangle = \sigma_2^2$; for simplicity we will assume that the noise is Gaussian. Intuitively, we should be able to do better by combining the two observations than we would do by looking just at one of them, and we also expect that we should give greater weight to the more accurate measurement. Quantitatively, if the measurements are independent of one another, we have

$$P(x|y_1, y_2) = \frac{P(y_1, y_2|x)P(x)}{P(y_1, y_2)} \quad (431)$$

$$\propto P(x)P(y_1|x)P(y_2|x) \quad (432)$$

$$\propto P(x) \exp \left[-\frac{1}{2\sigma_1^2}(y_1 - x)^2 - \frac{1}{2\sigma_2^2}(y_2 - x)^2 \right]. \quad (433)$$

Then we can form the optimal estimator in the least squares sense,

$$x_{\text{est}}(y_1, y_2) \equiv \int dx x P(x|y_1, y_2) \quad (434)$$

$$= \frac{\sigma_2^2 y_1 + \sigma_1^2 y_2}{\sigma_1^2 + \sigma_2^2}, \quad (435)$$

where in the last step we assume that the prior $P(x)$ is broad compared with the noise levels in our data. Thus, as expected, the optimal estimate is a combination of the data, and the weights are inverse to their relative noise levels.

Problem 76: Cue combination. Fill in the details leading to Eq (435). Can you work out the same problem but with additional multiplicative noise, $y_n = e^{g_n} x + \eta_n$, where g_n is also Gaussian? In this case, it is possible to generate errors that are very large, so presumably large disagreements between the data points y_1 and y_2 should not be resolved by simple averaging. See how much analytic progress you can make here, or do a simple simulation. This is deliberately open ended.

There are many situations in which we give strongly unequal weights to different data. A dramatic example is ventriloquism, in which we trust our eyes not our ears, and assign the source of speech to the person (or the dummy) whose lips are visibly moving. To see whether we are giving weights in relation to noise levels, as would be optimal, we have to do an experiment in which we can manipulate the effective noise levels. This was first done convincingly in tasks that require subjects to combine information from vision and touch, [add figure from Ernst & Banks, with explanation]. Although under normal conditions we give strong preference to our visual system, these data show convincingly that we do this only because our visual system provides much more accurate spatial information; if we can change their noise levels, people will change the weights given to different cues, as predicted by optimal estimation theory.

[loss functions, actions .. Maloney; Wolpert]

The examples of estimation that we have discussed thus far have in common that the distribution of the feature we are interested in estimating has a single well defined peak given the input sensory data. In many cases, however, the data that we collect with our senses have multiple interpretations, perhaps even multiple interpretations that provide equally good explanations of what we have seen or heard. These ‘ambiguous percepts’ arise in many contexts. When we experience these stimuli, our perceptions jump at random among the different possibilities. Could these random jumps originate from the same small noise sources that limit the reliability of our senses? [give fuller discussion, both visual and auditory examples ... alternative models ... maybe end with connection to conscious perception?]

Need to give a summary/conclusion for the section.

While there were many precursors, reaching back across centuries, the conclusive demonstration that bats navigate by echolocation, with sounds beyond the range of human hearing, was by Griffin & Galambos (1941). Griffin (1958) gives a beautiful presentation of the history and basic facts about the system. [need original ref for exp’t with dusted mealworms] The first suggestion of sub-microsecond precision in this system was from Simmons (1979).

Perhaps not surprisingly, these observations (and the provocative title of the paper in which they were presented) touched off a flurry of controversy; for different views, see Altes (1981) and Menne & Hackbarth (1986). The astonishing results on nanosecond precision, and the optimality of performance in background noise, were presented by Simmons et al (1990). For context, it is interesting to look at examples of precise timing measurements in binaural hearing [need ref, presumably to Konishi in barn owls] and in weakly electric fish (Rose & Heiligenberg 1985).

Altes 1981: Echo phase perception in bat sonar? RA Altes, *J Acoust Soc Am* **69**, 1232–1246 (1981).

Griffin 1958: *Listening in the Dark*. DR Griffin (Yale University Press, New Haven, 1958).

Griffin & Galambos 1941: The sensory basis of obstacle avoidance by flying bats. DR Griffin & R Galambos, *J Exp Zool* **86**, 481–506 (1941).

Menne & Hackbarth 1986: Accuracy of distance measurement in the bat *Eptesicus fuscus*: Theoretical aspects and computer simulations. D Menne & H Hackbarth, *J Acoust Soc Am* **79**, 386–397 (1986).

Rose & Heiligenberg 1985: Temporal hyperacuity in the electric sense of fish. G Rose & W Heiligenberg, *Nature* **318**, 178–180 (1985).

Simmons 1979: Perception of echo phase information in bat sonar. JA Simmons, *Science* **204**, 1336–1338 (1979).

Simmons et al 1990: Discrimination of jittered sonar echoes by the echolocating bat, *Eptesicus fuscus*: The shape of target images in echolocation. JA Simmons, M Ferragamo, CF Moss, SB Stevenson & RA Altes, *J Comp Physiol A* **167**, 589–616 (1990).

The program of comparing human performance with statistical limits in the context of higher level perception was outlined by Barlow (1980). The experiments on symmetry in random dot patterns are by Barlow & Reeves (1979), and an analysis of optimality in motion perception using random dot stimuli was given by Barlow & Tripathy (1997). For a review of how these stimuli have been used to probe the connections between neural activity and perception, see Newsome et al (1995). [Probably there needs to be a bit more here (!); maybe also in the text?] Note that, as discussed in *Spikes* (Rieke et al 1997; see below), these experiments connecting neural activity with perception in primates have been done, largely, in a regime where the subject is integrating imperfectly over very long periods of time, much longer than we would expect to see constant velocity motion in a natural setting; see also Osborne et al (2004). This complicates efforts to compare either neural or behavioral performance with the physical limits, and indeed I don't know of any effort to measure the responses of visual cortex in a regime (e.g., photon counting in the dark) where we understand fully the sources of noise limiting our perception; there is an opportunity here.

Barlow 1980: The absolute efficiency of perceptual decisions. HB Barlow, *Phil Trans R Soc Ser B* **290**, 71–82 (1980).

Barlow & Reeves 1979: The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. HB Barlow & BC Reeves, *Vision Res* **19**, 783–793 (1979).

Barlow & Tripathy 1997: Correspondence noise and signal pooling in the detection of coherent visual motion. H Barlow & SP Tripathy, *J Neurosci* **17**, 7954–7966 (1997).

Newsome et al 1995: Visual motion: Linking neuronal activity to psychophysical performance. WT Newsome, MN Shadlen, E Zohary, KH Britten & JA Movshon, in *The Cognitive Neurosciences*, M Gazzaniga, ed, pp 401–414 (MIT Press, Cambridge, 1995).

Osborne et al 2004: Time course of information about motion direction in visual area MT of macaque monkeys. LC Osborne, W Bialek & SG Lisberger, *J Neurosci* **24**, 3210–3222 (2004).

The classical work on motion estimation in insect vision was by Hassenstein and Reichardt (1956); perspectives on these early ideas are given by Reichardt (1961) and by Reichardt and Poggio (1976). A crucial piece of data in this discussion concerns the speed of a flying insect's motor response to visual motion, and a first estimate of this was given by Land and Collett (1974) in a beautiful analysis of natural flight trajectories; subsequent work was done by Wagner (1986a–c) and by Schilstra and van Hateren (1999; van Hateren & Schilstra 1999).

van Hateren & Schilstra 1999: Blowfly flight and optic flow. II. Head movements during flight. JH van Hateren & C Schilstra, *J Exp Biol* **202**, 1491–1500 (1999).

Hassenstein & Reichardt 1956: Systemtheoretische Analyse der Zeit-, Reihenfolgen-, und Vorzeichenbewertung bei der Bewegungserkennung des Rüsselkäfers. S Hassenstein & W Reichardt, *Z Naturforsch* **11b**, 513–524 (1956).

Land & Collett 1974: Chasing behavior of houseflies (*Fannia canicularis*): A description and analysis. MF Land & TS Collett, *J Comp Physiol* **89**, 331–357 (1974).

Reichardt 1961: Autocorrelation, a principle for the evaluation of sensory information by the central nervous system. W Reichardt, in *Sensory Communication*, WA Rosenblith, ed, pp 303–317 (MIT Press, Cambridge 1961).

Reichardt & Poggio 1976: Visual control of orientation behavior in flies. I. A quantitative analysis. W Reichardt & T Poggio, *Q Rev Biophys* **9**, 311–375 (1976).

Schilstra & van Hateren 1999: Blowfly flight and optic flow. I: Thorax kinematics and flight dynamics. C Schilstra & JH van Hateren, *J Exp Biol* **202**, 1481–1490 (1999).

Wagner 1986a: Flight performance and visual control of flight in the free-flying house fly (*Musca domestica* L.). I: Organization of the flight motor. H Wagner, *Phil Trans R Soc Lond Ser B* **312**, 527–551 (1986).

Wagner 1986b: Flight performance and visual control of flight in the free-flying house fly (*Musca domestica* L.). II: Pursuit of targets. H Wagner, *Phil Trans R Soc Lond Ser B* **312**, 553–579 (1986).

Wagner 1986c: Flight performance and visual control of flight in the free-flying house fly (*Musca domestica* L.). I: Interactions between angular movement induced by wide- and small-field stimuli. H Wagner, *Phil Trans R Soc Lond Ser B* **312**, 581–595 (1986).

[Introduce this with refs to the anatomy of the fly visual system.] Motion sensitive neurons in the fly visual system were discovered by Bishop & Keehn (1966), around the same time that Barlow et al (1964) discovered motion sensitive neurons in the rabbit retina. Today we take for granted that individual neurons can be selective for very complicated things, culminating in face- and object-selective neurons in the far reaches of the visual cortex [refs to Gross et al], but these early measurements were surprising. Indeed, in Barlow's hands, the observation of motion sensitivity played a key role in helping to shape the idea that cells respond to successively more complex conjunctions of features as we move through successive layers of processing [need to find which of HBB's refs is best here]. An early experiment showing that some of the motion sensitive neurons are a necessary link in optomotor behavior is by Hausen & Wehrhahn (1983); [since then ... ?].

Barlow et al 1964: Retinal ganglion cells responding selectively to direction and speed of image motion in the rabbit. HB Barlow, RM Hill & WR Levick, *J Physiol (Lond)* **173**, 377–407 (1964).

Bishop & Keehn 1966: Two types of neurones sensitive to motion in the optic lobe of the fly. LG Bishop & D G Keehn, *Nature* **212**, 1374–1376 (1966).

Hausen & Wehrhahn 1983: Microsurgical lesion of horizontal cells changes optomotor yaw responses in the blowfly *Calliphora erythrocephala*. K Hausen & C Wehrhahn, *Proc R Soc Lond B* **219**, 211–216 (1983).

The experiments on the precision of motion discrimination using the output of H1 are from de Ruyter van Steveninck & Bialek (1995), and the reconstruction of velocity waveforms was done in Bialek et al (1991); a review of these ideas and results is given in *Spikes* (Rieke et al 1997). A detailed calculation of the physical limits to motion estimation in this system is in my lecture notes from the Santa Fe Summer School (Bialek 1990). For a general discussion of hyperacuity in vision see Westheimer (1981), and for the relation of hyperacuity to physical limits, see Geisler (1984). The theory of optimal motion estimation is from Marc Potters' PhD thesis (Potters & Bialek 1994); related work was done by [need to understand exactly what Simoncelli and others did around the same time], and application of these ideas to human visual motion perception can be found in Weiss et al (2002). Problem [**] about third order statistics is inspired by Fitzgerald et al (2011). [Stocker?]

Bialek 1990: Theoretical physics meets experimental neurobiology. W Bialek, in *1989 Lectures in Complex Systems, SFI Studies in the Sciences of Complexity, Vol II*, E Jen, ed, pp 513–595 (Addison–Wesley, Menlo Park CA, 1990).

Bialek et al 1991: Reading a neural code. W Bialek, F Rieke, RR de Ruyter van Steveninck & D Warland, *Science* **252**, 1854–1857 (1991).

Fitzgerald et al 2011: Symmetries in stimulus statistics shape the form of visual motion estimators. JE Fitzgerald, AY Katsov, TR Clandinin & MJ Schnitzer, *Proc Nat'l Acad Sci (USA)* in press (2011).

Geisler 1984: Physical limits of acuity and hyperacuity. WS Geisler, *J Opt Soc Am A* **1**, 775–782 (1994).

Potters & Bialek 1994: Statistical mechanics and visual signal processing. M Potters & W Bialek, *J Phys I France* **4**, 1755–1775 (1994).

Rieke et al 1997: *Spikes: Exploring the Neural Code*. F Rieke, D Warland, R de Ruyter van Steveninck & W Bialek (MIT Press, Cambridge, 1997).

de Ruyter van Steveninck & Bialek 1995: Reliability and statistical efficiency of a blowfly movement-sensitive neuron. R de Ruyter van Steveninck & W Bialek, *Phil Trans R. Soc Lond Ser B* **348**, 321–340 (1995).

Weiss et al 2002: Motion illusions as optimal percepts. Y Weiss, EP Simoncelli & EH Adelson, *Nature Neurosci* **5**, 598–604 (2002).

Westheimer 1981: Visual hyperacuity. G Westheimer, *Prog Sens Physiol* **1**, 1–30 (1981).

The classical evidence for the systematic errors of motion estimation predicted by the correlator model are discussed by Reichardt & Poggio (1976), above. Experiments showing the quadratic contrast dependence of responses in the motion sensitive neurons include [need to find the early ones!]. The demonstration that quadratic behavior at low contrasts coexists with unambiguous responses to velocity at high contrast is given by de Ruyter van Steveninck et al (1994, 1996) [check that these are the best references!]. These experiments were done with randomly textured images, whereas classical studies of visual motion have used periodic gratings. The correlator model also predicts that velocity will be confounded with the spatial frequency of these gratings, and this error persists even under high signal-to-noise ratio conditions (Haag et al 2004); it is not clear whether this represents a genuine failure of optimal estimation, a byproduct of strategies for gain control and efficient coding (Borst 2007), or simply a behavior that would never be seen under natural conditions. There are several experiments, especially in bees (Srinivasan et al 1991, 1996; Baird et al 2005),

indicating that insects have access to signals that allow them to control their flight speed without any of the systematic errors predicted by the correlator model; recent work confirms this conclusion in *Drosophila* using sophisticated tracking and virtual reality to allow control experiments under free flight conditions (Fry et al 2009). A number of experiments have shown that the responses of motion-sensitive neurons are also very different under more natural conditions (Lewen et al 2001, de Ruyter van Steveninck et al 2001), although most of the analysis has focused on the nature of coding in spike trains rather than the nature of the motion computation itself. [Is Rob's experiment on reducing ambiguity at high light levels, outside, published?] An attempt to dissect the motion computation represented by the spiking output of H1 is described in Bialek & de Ruyter van Steveninck (2005), and in Appendix A.7. [Do we say something about controversies?]

Baird et al 2005: Visual control of flight speed in honeybees. E Baird, MV Srinivasan, S Zhang & A Cowling, *J Exp Biol* **208**, 3895–3905 (2005).

Bialek & de Ruyter van Steveninck 2005: Features and dimensions: Motion estimation in fly vision. W Bialek & R de Ruyter van Steveninck, arXiv:q-bio/0505003 (2005).

Borst 2007: Correlation versus gradient type motion detectors: the pros and cons. A Borst, *Phil Trans R Soc Lond Ser B* **362**, 369–374 (2005).

Fry et al 2009: Visual control of flight speed in *Drosophila melanogaster*. SN Fry, N Rohrseitz, AD Straw & MH Dickinson, *J Exp Biol* **212**, 1120–1130 (2009).

Lewen et al 2001: Neural coding of naturalistic motion stimuli. GD Lewen, W Bialek & RR de Ruyter van Steveninck, *Network* **12**, 317–329 (2001); arXiv:physics/0103088 (2001).

de Ruyter van Steveninck et al 1994: Statistical adaptation and optimal estimation in movement computation by the blowfly visual system. RR de Ruyter van Steveninck, W Bialek, M Potters & RH Carlson, in *Proc IEEE Conf Sys Man Cybern*, 302–307 (1994).

de Ruyter van Steveninck et al 1996: Adaptive movement computation by the blowfly visual system. RR de Ruyter van Steveninck, W Bialek, M Potters, RH Carlson & GD Lewen in *Natural and Artificial Parallel Computation: Proceedings of the Fifth NEC Research Symposium*, DL Waltz, ed, 21–41 (SIAM, Philadelphia, 1996).

de Ruyter van Steveninck et al 2001: Real time encoding of motion: Answerable questions and questionable answers from the fly's visual system. R de Ruyter van Steveninck, A Borst & W Bialek, in *Processing Visual Motion in the Real World: A Survey of Computational, Neural and Ecological Constraints*, JM Zanker & J Zeil, eds, pp 279–306 (Springer-Verlag, Berlin, 2001); arXiv:physics/0004060 (2000).

Srinivasan et al 1991: Range perception through apparent image speed in freely flying honeybees. MV Srinivasan, M Lehrer, WH Kirchner & SW Zhang, *Vis Neurosci* **6**, 519–535 (1991).

Srinivasan et al 1996: Honeybee navigation en route to the goal: visual flight control and odometry. MV Srinivasan, S Zhang, M Lehrer & TS Collett, *J Exp Biol* **199**, 237–244 (1996).

Since that formative year of having the office next door to Rob de Ruyter van Steveninck when I was a postdoc in Groningen, the fly visual system has seemed to me an ideal testing ground for physicists' ideas. On the other hand, if you think that brains are interesting because you want to understand your own brain, you might believe that insects are a bit of a side show relative to animals that share more of our brain structures—monkeys, cats, or even mice. There are obvious questions of strategy here, including the fact that (perhaps paradoxically) it can be easier to control the behavior of a primate than the behavior of an insect, creating

opportunities for certain kinds of quantitative experiments. There also are questions about how much universality we should expect. Are there things to be learned about brains in general, or is everything about our brain different from that of “lower” animals? Can careful, quantitative analyses of “simpler” systems sharpen the questions that we ask about bigger brains (even if the answers are different), or does each case present such unique challenges? I think it is fair to say that for several decades there has been a strong consensus of the mainstream neuroscience community that the answers to these questions point away from the study of insect brains. Recently, however, there has been substantial growth in a community of scientists interested in exploiting the tools of modern molecular biology to study the brain, and this group of course is attracted to “model organisms” with well developed methods of genetic manipulation, such as the fruit fly *Drosophila melanogaster* and its close relatives. Thus, the coming years are likely to see a resurgence of interest in insect brains, and this should create more opportunities for physicists. It is early days, but here is a selection of papers that may help you in your explorations.

: Find a selection of *Drosophila* articles that point toward quantitative opportunities.

Seelig et al 2010: Two-photon calcium imaging from head-fixed *Drosophila* during optomotor walking behavior. JD Seelig, ME Chiappe, GK Lott, A Dutta, JE Osborne, MB Reiser & V Jayaraman, *Nature Methods* **7**, 535–540 (2010).

The rather astonishing results in Fig 67 are from Griffiths & Tennenbaum (2006). The original work on optimal cue combination was by Ernst & Banks (2002) [cite follow ups!]. [More: Maloney, Wolpert, Finally, ambiguous percepts, multistability, connections to conscious experience ...]

Bialek & DeWeese 1995: Random switching and optimal processing in the perception of ambiguous signals. W Bialek & M DeWeese, *Phys Rev Lett* **74**, 3077–3080 (1995).

Ernst & Banks 2002: Humans integrate visual and haptic information in a statistically optimal fashion. MO Ernst & MS Banks, *Nature* **415**, 429–433 (2002).

Griffiths & Tennenbaum 2006: Optimal predictions in everyday cognition. TL Griffiths & JB Tennenbaum, *Psychological Science* **17**, 767–773 (2006).

D. Proofreading and active noise reduction

Fluctuations are an essential part of being at thermal equilibrium. Thus, the fact that life operates in a relatively narrow range of temperatures around 300 K means that some level of noise is inevitable. But being alive certainly is not being at thermal equilibrium. Can organisms use their non-equilibrium state to reduce the impact of nominally thermal noise? More generally, can we understand how to take a system in contact with an environment at temperature T , and expend energy, driving it away from equilibrium, in such a way as to reduce the effects of noise?

In his classic lectures *What is Life?*, Schrödinger waxed eloquent about the fidelity with which genetic information is passed from generation to the next, conjuring the image of a gallery with portraits of the Hapsburgs, their

oddly shaped lips reproduced across centuries of descendants. Schrödinger was much impressed by the work of Timoféef-Ressovsky, Zimmer and Delbrück, who had determined the cross-section for ionizing radiation to generate mutations, and used this to argue that genes were of the dimensions of single molecules. Thus, the extreme stability of our genetic inheritance could not be based on averaging over many molecules, as a “naive classical physicist” might have thought. Now is a good time to set aside our modern insouciance and allow our ourselves to be astonished, as Schrödinger was, that so many of the phenomena of life are the macroscopic consequences of individual molecular events.

We now teach high school students that the key to the transmission of genetic information is the pairing of bases along the double helix—A pairs with T, C pairs with G, as in Fig 68. This, of course, is the triumph of Watson and Crick’s theory of DNA structure.⁵² The ideas of templates and structural complementarity which are at the heart of the double helix reappear many times—every time, in fact, that the organism needs to make reliable choices about which molecules to synthesize. But does structural complementarity solve the problem of reliabil-




FIG. 68 Base pairing in the Watson-Crick structure of DNA [this just grabbed from Wikipedia; need to decide exactly what to show, and redraw]. At left, we see the hydrogen bonding between bases in the correct pairings, showing how they “fit” to satisfy the opportunities for hydrogen bonding, producing structures that are the same width and hence can fit into the double helix, as shown at right. “R” denotes the sugar and phosphate groups, identical for all bases, which form the outer backbone(s) of the helix.

⁵² It would be almost silly to think you know something about “biophysics” (whatever you think the word means!) and not understand the interplay of theory and experiment that led to this revolution in the middle of the twentieth century. For a brief tour, see Appendix A.5.

ity in biosynthesis?

The fact that A pairs with T is really the statement that the (free) energy of a correct AT pair is much lower than that of an incorrect AC or AG pair. We should recall the energy scales for chemical bonding. A genuine covalent bond, such as the carbon–carbon or carbon–nitrogen bonds in the interior of the bases, results from the sharing of electrons between the atoms, and the energies are therefore on the scale of several electron volts.⁵³ Making the wrong base pairs wouldn't require us to break any covalent bonds, so the energy cost will not be this large. If we tried to make an AG pair, it would be so big that it wouldn't fit inside the backbone of the double helix; more precisely, we would have to make large distortions of the covalent bonds, and since these are stiff, the energy cost would be very large. On the other hand, if we try to make a CT pair, the backbone will hold the bases so far apart that they can't form hydrogen bonds. Thus, the minimal energy for a “wrong” base pair is the energy of two missing hydrogen bonds, and this is on the order of $10 k_B T$.

An energy difference of $\Delta F \sim 10 k_B T$ means that the probability of an *incorrect* base pairing should be, according to the Boltzmann distribution, $e^{-\Delta F/k_B T} \sim 10^{-4}$. A typical protein is three hundred amino acids long, which means that is encoded by nearly one thousand bases; if the error probability is 10^{-4} , then replication of the DNA would introduce roughly one mutation in every tenth protein. For humans, with a billion base pairs in the genome, every child would be born with hundreds of thousands of bases different from his or her parents. If these predicted error rates seem large, they are—real error rates in DNA replication vary across organisms, but are in the range of $10^{-8} - 10^{-12}$, so that entire genomes can be copied almost without any mistakes.

The discrepancy between Boltzmann probabilities and observed error rates is much more widespread. When information encoded in the DNA is read out to make proteins, there are several steps where errors can occur. First is the synthesis of mRNA from the DNA template, a process not unlike the replication of the DNA itself. The “codebook” for translating from the language of bases along the mRNA into amino acids is embodied in the tRNA molecules, which at one end have a triplet of bases (the anti-codon) that is complementary to a particular triplet of bases along the mRNA (the codon), and at their other end is the amino acid that the codon represents. To make such molecules, there are specialized enzymes that recognize the ‘bare’ tRNA and choose out of the cellular

soup the correct amino acid with which to ‘charge’ the molecule. [the discussion of tRNA and charging could use some sketches!] But some amino acids differ simply by the replacement of a CH_3 group with an H; if we imagine the enzyme recognizing the first amino acid with a binding pocket that is complementary to the CH_3 group, then the second amino acid will also fit, and the binding energy will be weaker only by the loss of non-covalent contacts with the methyl group; it is difficult to see how this could be much more than $\sim 5 k_B T$, corresponding to error rates $\sim 10^{-2}$. If the error rates in tRNA charging were typically 10^{-2} , almost all proteins would have at least one wrong amino acid; in fact error rates are more like 10^{-4} , so that most proteins have no errors. There is one more step, at the ribosome, where tRNA molecules bind to their complementary sites along the mRNA and the amino acids which they carry are stitched together into proteins, and here too there is a discrepancy between thermodynamics and the observed error probabilities.

Each of the events we have outlined—DNA replication, mRNA synthesis, tRNA charging, and protein synthesis on the ribosome—has its own bewildering array of biochemical details, and is the subject of its own vast literature. As physicists we search for common theoretical principles that can organize this biological complexity, and I think that this problem of accuracy beyond the thermodynamic limit provides a wonderful model for this search. The key ideas go back to Hopfield and Ninio in the 1970s. Their classic papers usually are remembered for having contributed to the solution of the problem of accuracy, a solution termed ‘kinetic proofreading,’ which we will explore in a moment. But I think they should also be remembered for having recognized that there is a common physics problem that runs through this broad range of different biochemical processes.

To understand the essence of kinetic proofreading, it is useful to recall the problem of Maxwell's demon. Imagine a container partitioned into two chambers by a wall, with a small door in the wall. [again, a sketch would help!] Maxwell conjured the image of a small demon who controls the door. If he⁵⁴ see a molecule coming from the right at high speed, he opens the door and allows it to go into the left chamber. Conversely, if he sees a molecule drifting slowly from the left, he opens the door and allows it to enter the right chamber. After some time, all the slow molecules are on the right, all the fast molecules are on the left. But, since the average kinetic energy of the molecules in a gas is proportional to the temperature, the demon has created a temperature difference, hot on the left, cold on the right. This temperature difference can be used to do useful work (e.g., running a heat engine), and thus the demon appears to have created something out

⁵³ Chemists prefer to think per mole rather than per molecule, and they prefer joules to electron Volts (I won't speak of calories). To have some numbers at your fingertips, remember that at room temperature, $k_B T = 1/40 \text{ eV} = 2.5 \text{ kJ/mole}$.

⁵⁴ Why is it obvious that the demon is male?

of nothing, violating the second law of thermodynamics.

There is nothing special about the demon's choice of molecular speed as the criterion for opening the door. It is a simple choice, because the result is a temperature difference, and we can imagine all sorts of appropriately nineteenth century methods for extracting useful work from temperature differences. But if there are two kinds of molecules, A and B , and the demon arranges for the A molecules to accumulate in the left chamber and B molecules to accumulate in the right chamber, then there will be differences in chemical potential between the two chambers, and there must be some way of using this to do work even if as physicists we don't know enough chemistry to figure it out.

Problem 77: Pushing away from equilibrium. Consider a polymer made from A and B monomers. Suppose we start with pure poly- A , and use this as a template to construct a new polymer, much as in DNA replication (but simpler!). Template directed synthesis works because the $A-A$ bond is stronger than the $A-B$ bond by some free energy difference ΔG ; we'll use the convention that $\Delta G > 0$. Then if we make a polymer of length N in which a fraction f of the monomers are incorrectly made to be B rather than A , the free energy of the system will have a contribution $Nf\Delta G$ relative to the perfectly copied poly- A . If the errors are made at random, however, then there is a contribution to the entropy of the polymer that comes from the sequence heterogeneity.

(a.) Evaluate the entropy that comes from the random substitutions of A by B . What assumptions are you making in this calculation? Can you imagine these being violated by real molecules?

(b.) Combine the entropy from [a.] with the "bonding" free energy $Nf\Delta G$ to give the total free energy of the polymer. Show that this is minimized at $f_{\text{eq}} \propto \exp(-\Delta G/k_B T)$, as expected.

(c.) How much free energy is stored in the polymer when $f < f_{\text{eq}}$? Can you give simple expressions when the difference $f_{\text{eq}} - f$ is small? What happens if (as we will see below) $f \approx f_{\text{eq}}^2$?

The demon's sin is to have generated a state of reduced entropy. We know that to enforce the second law, this non-equilibrium state must be 'paid for' with enough energy to balance the books—to avoid building a perpetual motion machine, the demon must have dissipated an amount of energy equal to or greater than the amount of useful work that can be extracted from his reduction in the entropy of the system. The key insight of Hopfield and Ninio was that the problem of accuracy or low error rates was of this same kind: achieving low error rates, sorting molecular components with a precision beyond that predicted by the Boltzmann distribution, means that the cell is building and maintaining a non-equilibrium state, and it must spend energy in order to do this. Somewhere in the complexity of the biochemistry of these processes there must be steps which dissipate energy, and this has to be harnessed to improve the accuracy of synthesis.




FIG. 69 The simplest kinetic scheme in which an enzyme can choose correct or incorrect molecules out of solution, making correct or incorrect products.

To see how this might work, let's look at the simplest model of a biochemical process catalyzed by an enzyme, as in Fig 69. In essence, the chemical reaction of interest involves choosing among two (or more) substrate molecules, for example the correct and incorrect base at a particular point along the strand of DNA that the cell is trying to replicate or transcribe into mRNA. In order to complete the reaction, the substrate has to bind to the enzyme, and this enzyme-substrate complex can be converted into the product; in order to have any possibility of correcting errors, it must be possible for the substrate to unbind from the enzyme before the conversion to product. With only this minimum number of steps, the kinetics are described by

$$\frac{d[EA]}{dt} = k_+[A][E] - (k_- + V_{\max})[EA] \quad (436)$$

$$\frac{d[EB]}{dt} = k'_+[B][E] - (k'_- + V'_{\max})[EB] \quad (437)$$

$$[E]_{\text{total}} = [EA] + [EB] + [E], \quad (438)$$

where A is the correct substrate, B is the incorrect substrate, and $[E]_{\text{total}}$ is the (fixed) total concentration of enzyme molecules. The rate at which correct products are made is given by $V_{\max}[EA]$, and the rate of making incorrect products is $V'_{\max}[EB]$. If the overall rate of reactions is slow enough not to deplete the substrates (and the cell typically is working hard to make sure this is true!), then we can compute these rates in the steady state approximation.

To compute the rate of errors we don't even need to solve the entire problem. From Eq (436) we can see that, in steady state,

$$[EA] = [E] \frac{k_+[A]}{k_- + V_{\max}}; \quad (439)$$

similarly, from Eq (437),

$$[EB] = [E] \frac{k'_+[B]}{k'_- + V'_{\max}}. \quad (440)$$

Thus the error probability, or relative rate at which incorrect products are made, is given by

$$f \equiv \frac{\text{rate of making incorrect product}}{\text{rate of making correct product}} \quad (441)$$

$$= \frac{V'_{\max}[EB]}{V_{\max}[EA]} \quad (442)$$

$$= \left[\frac{k'_+[B]}{k'_- + V'_{\max}} \right] \times \left[\frac{k_+[A]}{k_- + V_{\max}} \right]^{-1} \times \left[\frac{V'_{\max}}{V_{\max}} \right] \quad (443)$$

To go further it is useful to notice that all the reactions we are thinking about share one important feature: the actual making and breaking of covalent bonds occurs on ‘the other side’ of the molecule from the structure that defines correct vs. incorrect [definitely needs a sketch!]. In the case of DNA replication, for example, correctness has to do with the pattern of hydrogen bonding between the bases, on the inside of the helix, while the actual reaction required to incorporate one base into the growing polymer involves the phosphate backbone on the outside of the helix. This makes it unlikely that the rate at which these bonds are formed is sensitive to the correctness of the substrate. Correspondingly, in the cases of interest, it is likely that $V'_{\max} \approx V_{\max}$, so this is not a source of selectivity. More importantly, from Eq (443) it is clear that, under these conditions, the error probability is minimized if the catalytic rate V_{\max} is slow compared with the unbinding rates k_-, k'_- . This makes sense: if the catalytic step itself has no selectivity, then to maximize selectivity one must give the wrong substrate a chance to fall off.

So, when the dust settles, in this simplest kinetic scheme we have shown that the error probability is bounded,

$$f > \left(\frac{k'_+[B]}{k'_-} \right) / \left(\frac{k_+[A]}{k_-} \right). \quad (444)$$

But this combination of rates and concentrations is exactly what determines the equilibrium binding of A vs B to the enzyme, and hence can be written in terms of thermodynamic quantities,

$$f > \exp \left(-\frac{F_A - F_B}{k_B T} \right), \quad (445)$$

where F_A is the free energy for taking a single molecule of A out of solution and binding to the enzyme, and similarly for B ; here binding energies are positive, larger for tighter binding. Thus, we are back where we started, with an error probability determined by the Boltzmann distribution!




FIG. 70 The simplest scheme for “kinetic proofreading.” As described in the text, the key step is an irreversible transition from EA to EA^* , which gives a true second chance for equilibration with the free A molecules.

But the Michaelis–Menten scheme has a natural generalization. Suppose that, after binding, there is an irreversible transition to a new state, at a rate r , and that in this state the substrate can again be released from the enzyme, as in Fig 70. In the simplest case, the events which determine binding and release of the (perhaps modified) substrate are the same as in the initial step, with the same rates. We can carry through the analysis of this kinetic scheme as before, and with the same assumption that catalytic steps (V_{\max} and r) have no selectivity, find that

$$f > \exp \left(-\frac{F_A - F_B}{k_B T} \right) \exp \left(-\frac{F_{A^*} - F_{B^*}}{k_B T} \right). \quad (446)$$

But if the molecular interactions that select A over B are the same for A^* vs B^* , we expect $F_{A^*} - F_{B^*} \approx F_A - F_B$, and hence

$$f \rightarrow \left[\exp \left(-\frac{F_A - F_B}{k_B T} \right) \right]^2. \quad (447)$$

This is the essence of kinetic proofreading: by introducing an irreversible step into the kinetic scheme, a step which necessarily dissipates energy, it is possible to use the equilibrium selectivity twice, and achieve an error probability which is the square of the nominal limit set by the Boltzmann distribution.

Problem 78: More on the basics of kinetic proofreading.

To begin, give the details needed to derive Eq (446). An even better exercise is to go through Hopfield’s original paper (Hopfield 1974), pen in hand, filling in all the missing steps. Then consider the following:

(a.) In the simplest scheme, we saw that maximum selectivity occurs when V_{\max} is slow compared with k_- . Is there a similar condition in the proofreading scheme? What does this tell us about the progress of the enzymatic cycle? More specifically, what is the fate of the typical substrate which binds to the enzyme? Is it converted to product, or ejected as A^* ?

(b.) Consider a generalization of the kinetic scheme in Fig 70 such that the nominally irreversible step with rate r is in fact reversible, with the reverse reaction at rate r' . To be general, imagine also the binding and unbinding of A^* can occur with rates that are different from the rates for A . Now there are detailed balance conditions that connect these different rates. Write down these conditions, and show how they effect the error probability. Can you say something general here? In particular, can you show how these conditions enforce the Boltzmann error rate in the absence of energy dissipation, no matter how many times the enzyme ‘looks’ at the substrate?

How does this general idea of proofreading connect with the real biochemistry of these systems? In some sense the case of DNA replication (or transcription) is most obvious, as shown in Fig 71. All of the nucleotides which are incorporated into the growing strands of DNA or RNA start as nucleotide triphosphates, but once the final structure is formed only one phosphate is part of the backbone. Thus, at some point in the process, the ‘high energy’ phosphate bond must be cleaved, releasing roughly $20k_B T$ of free energy. If this is the irreversible step, then it must be possible for the enzyme which catalyzes the growth of the polymer to release the nucleotide after this cleavage, which means after it has been attached to the backbone of the growing chain. Thus, to




FIG. 71 Connecting the proofreading scheme to specific biochemical process, from Hopfield (1974). At the top, nucleotide triphosphates are incorporated as monophosphates in DNA replication or the transcription to mRNA. In the middle panel, the charging of tRNA molecules with amino acids, involving an extra ATP. At bottom, a very simplified view of protein synthesis, in which the GTP/exchange by the protein Tu provides the energy for proofreading at the ribosome.

proofread, the enzyme must be not only a ‘polymerase’ (catalyzing the polymerization reaction) it must also be an ‘exonuclease’ (catalyzing the removal of nucleotides from the polymer). It had been known almost since the discovery of the polymerase that it also had exonuclease activity, but it took the idea of kinetic proofreading to explain how this was connected, through energy dissipation, to proofreading and error correction. In the charging of tRNA, the process actually starts with an ATP molecule being cleaved, leaving an AMP attached to the amino acid before it reacts with the tRNA. In protein synthesis, the sequence of reactions is much more complex, but again there is an obligatory cleavage of a nucleotide triphosphate (in this case $\text{GTP} \rightarrow \text{GDP}$). All of these examples are qualitatively consistent with the proofreading scenario,⁵⁵ and especially in the case of tRNA charging it has been possible to pursue a more quantitative connection between theory and experiment [do we want to say more about this?].

Kinetic proofreading not only solves a fundamental problem—the problem which Schrödinger confronted in the Hapsburg portraits—it also has been a source of new questions and ideas. If the accuracy of DNA replication depends not only on intrinsic properties of the DNA but also on the detailed kinetics of the enzymes involved in replication, then the rate of mutations itself can be changed by mutations. It has long been known that there are ‘mutator strains’ of bacteria which have unusually high error rates, and we now know that that these strains simply have aspects of the proofreading apparatus disabled. One could imagine subtler changes, so that the mutation rate would become a quantitative trait; in this case the dynamics of evolution would be very different, since fluctuations along one “direction” in the space of genomes would change the rate of movement along all directions. Also, since accuracy depends on energy dissipation, in an environment with limited nutrients there is a tradeoff between the speed of growth and the fidelity with which genetic information is passed to the next generations; there is an optimization problem to be solved here, and ... [say something definite re Kurland, Ehrenberg, ... maybe have a problem?]. In protein synthesis, accuracy and even the overall kinetics will be affected by the availability of the different charged tRNAs, and this is under physiological control, so again there is the pos-

⁵⁵ Hopfield has also emphasized that there are kinetic schemes in which proofreading still proceeds through energy dissipating steps, but if the enzymes have some memory for past events then the synthesis and dissipation can be separated in time, erasing some of the more obvious signatures from the simpler scheme. This may be especially important in thinking about more complex examples, such as protein synthesis on the ribosome or DNA replication in higher eukaryotes. [is there a good problem to give here?]

sibility that, especially for fast growing bacteria where the problems are most serious, there is some tuning or optimization to be done.

Problem 79: Controlling the pace of evolution? [take the students through a simple version of Magnasco & Thaler. Introduces ideas of evolutionary landscape, connect back to discussion of reaction rates ...]

Problem 80: Optimizing tRNA pools. There is a separate tRNA complementary to each of the 60 codons which code for amino acids (the remaining four codons stand for ‘start’ and ‘stop’). The frequency with which these codons are used in the genome varies widely, both because proteins do not use all 20 amino acids equally and because different organisms use different synonymous codons (that is, those which code for the same amino acid) with different frequencies. But, when it comes time to make protein, the cell needs access to the appropriate population of charged tRNAs. Naively one might expect that, if the supply of tRNA is limiting the rate at which a bacterium can make proteins and grow, then it would be good to have a supply of tRNA in proportion to how often the corresponding codon gets used. Let’s see if this is right. Suppose that protein synthesis is limited by arrival of the tRNA at the ribosome. Then the time required to incorporate one amino acid coded by codon i is $t_i \sim 1/k[tRNA_i]$, where k is a second order rate constant.

(a.) [try to sort out how rate of ribosome turnover compares with diffusion limited rate of arrival of tRNAs]

(b.) The average time required to incorporate one amino acid is $\bar{t} = \sum_i p_i/k[tRNA_i]$, where p_i is the probability of codon i appearing in the cell’s mRNA. If the cell can only afford a limited amount of tRNA, the natural constraint is on the total $\sum_i [tRNA_i]$. How should the individual concentrations be arranged to minimize the mean incorporation time \bar{t} ? Is this surprising?

(c.) You might be tempted to say that, if the goal is to synthesize proteins as rapidly as possible, and the rates are limited by the arrival of tRNAs, then we should maximize the mean rate, $\sum_i p_i k[tRNA_i]$. Why is this wrong?

The ideas of kinetic proofreading may be even more generally applicable than envisioned by Hopfield and Ninio. There are many signal transduction processes that start with a receptor binding event at the cell surface and trigger a cascade of protein phosphorylation reactions;⁵⁶ the phosphate groups are pulled from ATP, so phosphorylation is a prototypically irreversible, energy consuming reaction. In the immune system [need a figure here!] it has been suggested that this can provide multiple stages

⁵⁶ [Should have said something about this already!] Many proteins are activated by the covalent addition of phosphate groups, a reaction termed phosphorylation. Enzymes that catalyze the transfer of phosphate groups are called kinases, and these enzymes often are usually specific for their substrates, whether these are smaller molecules or proteins. Importantly, some kinases themselves are activated by phosphorylation, and the enzymes that carry out this activation step are termed kinase kinases.




FIG. 72 Kinetic proofreading in the phosphorylation of a kinase (K) by a kinase-kinase (KK), from Swain & Siggia (2002). Activation of the kinase requires two steps of phosphorylation, and in this scheme the kinase-kinase can dissociate from its substrate after have transferred just one phosphate group. K_0 , K_1 and K_2 denote the kinase with zero, one and two attached phosphate groups, respectively.

of proofreading, contributing to self/non-self discrimination. More generally, as shown in Fig 72, if activation of an enzyme requires two steps of phosphorylation, then these steps can be arranged in a proofreading scheme. Because there are many such pathways in the cell, proofreading this case could increase specificity and reduce crosstalk.

Watson and Crick understood that the double helical structure of DNA, with its complementary strands, suggested a mechanism for the copying of genetic information from one generation to the next. But they also realized that the helical structure creates a problem, since the strands are entangled; the problem is most obvious in bacteria, where the chromosomes close into circles, but with very long molecules one couldn’t rely on spontaneous untying even if there is no formal topological obstruction. Eventually it was discovered that there is a remarkable set of enzymes that catalyze changes in the topology of circular DNA molecules, allowing the strands to pass through one another. In the process of relieving entanglement, these “topoisomerases” also reduce the energy stored in the supercoiling of these polymers [should say more about this here—an excuse to talk about link, writhe and twist, etc.; certainly needs a figure]. The problem is that being truly unlinked is a global property of the molecules, while the enzymes act locally. In the simplest models, then, topoisomerases would remove the obstacles to changing topology, but couldn’t shift the probability of being unlinked from its equilibrium value. Because making links or knots restricts the entropy of the molecule, there is an equilibrium bias in favor of un-linking, but this seems insufficient for cellular function.

Indeed, as shown in Fig 73, topoisomerases seem to leave fewer links than expected from the Boltzmann distribution even in test tube experiments, and if we look at the details of the biochemical steps involved, we can identify a series of steps that are equivalent to proofreading by the topoisomerases [I'd like to explain this better!].

The ideas of proofreading have recently been revitalized by the opportunity to observe, more or less directly, the individual molecular events responsible for error correction. The key to this new generation of experiments is the realization that molecules such as RNA polymerase are “molecular motors” that move along the DNA strand as they function. Each step in this movement is presumably on the scale of the distance between bases along the DNA, $d \sim 3.4 \text{ \AA}$. The energy to drive this motion comes from breaking the phosphate bonds of the input nucleotides, and is on the scale of $\sim 10k_B T$. Thus the forces involved are $F \sim 10k_B T/d \sim 100 \text{ pN}$.

When a dielectric sphere sits in an electric field, it polarizes, and the direction of the polarization is such that it lowers the energy. This means that the energy of the sphere is lower in regions of high electric field. Since the energy is proportional to the square of the field, this is true even if the field is oscillating in time. In particular, if we focus a light beam in a microscope, then the light intensity is higher in the focus, and light intensity is just the square of the electric field, so we expect that small dielectric objects will be attracted to focal spots, and this




FIG. 73 Kinetic proofreading in DNA unlinking, from Yan et al (1999). At left, experimental results redrawn from Rybenkov et al (1997), showing that topoisomerases reaching a linking probability roughly equal to the square of the expected equilibrium probability, suggesting a proofreading scheme. At right, a kinetic scheme illustrating the possibility of proofreading. Active topoisomerase molecules are shown in red, inactive in blue; green arrows denote transitions that are insensitive to the topology, while all sensitivity is contained in the red arrows. This kinetic scheme is essentially a “folded” version of Hopfield’s original Fig 70.




FIG. 74 Schematic of an experiment to observe the function of RNA polymerase with single base-pair resolution, from Shaevitz et al (2003). A laser beam is split, and the two resulting beams are focused to make “optical traps” for two micron-sized beads. Attached to one bead is a double stranded DNA molecule, and attached to the other is an RNA polymerase molecule. As the polymerase synthesizes mRNA, it “walks” along the DNA and the tether between the two beads is shortened. The intensities of the two beams are set so that the left hand trap is stiffer, insuring that most of the motion appears as a displacement of the right hand bead, which is measured by projecting scattered light onto a position-sensitive detector.

is called “optical trapping.” Importantly, with realistic light intensities, the forces on micron-sized particles as they move in an optical trap indeed are on the scale of piconewtons, so it is possible to “hold” a molecular motor in place.

Problem 81: Optical trapping. The key to the experiments here is the fact that small, neutral particles can be trapped at the focus of a laser beam, and that the forces generated in this way are on the same scale as those generated by individual biological motor molecules, such as the RNA polymerase. **Take the students through this!**

In Figure 74 we see the schematic of an optical trapping experiment on the RNA polymerase. Successive generations of technical improvements in these experiments have made it possible to track the motion of the polymerase with a resolution fine enough to see it “step” from one base pair to the next, as in Fig 75. Importantly, in these experiments one can bathe the sample in a solution containing different nucleotides. If we add ITP, which is not one of the standard four bases, it will sometimes be incorporated into the growing mRNA strand, but this is




FIG. 75 Motion of the RNA polymerase along DNA. At left, from Abbondanzieri et al (2005). Top, the position of the right hand bead from Fig 74 as the trap is moved in 1 Å steps, to show that these can be resolved. Bottom, the active motion of the bead as the RNA polymerase synthesizes rRNA, showing the expected steps of 3.4 Å. [are the black lines median filtering?] At right, from Shaevitz et al (2005). Top, the average trajectory of the RNA polymerase aligned on the start and end of long pauses. Bottom, the mean duration of pauses under different conditions, notably the addition of the “wrong” nucleotide ITP.

always a mistake. Under these conditions we can observe an increased frequency of “pauses” in the motion of the polymerase, followed by backtracking of 1–10 base pairs along a relatively stereotyped trajectory. If we remove from RNA polymerase the subunits thought to be involved in proofreading, then these error-induced pauses become very long.

[Need a summary on kinetic proofreading, segue to active filtering]

There is another broad class of examples in which there seems to be a discrepancy between the noise expected at thermal equilibrium and the performance of biological systems, and this is in the measurement of small displacements. In our inner ear, and in the ears of all other vertebrate animals, motions are sensed by “hair cells,” so named because of the tuft of “hairs” (more properly, stereocilia) that project from their top surface as in Fig 76. Although we usually think of ears as responding to airborne sounds, in fact there are multiple chambers in the ear, some of which respond to sound, and others of which respond to lower frequency motions generated by rotation of our head, the largely constant force of gravity or ground borne vibrations. The core of all these systems, however, is the hair cell. When the stereocilia are bent, channels in the cell membrane open and close, and this modulates an ionic current, as in other receptor cells that we have seen before. In a variety of systems it has been possible to open these organs, or

even dissect out the hair cells, and to make direct mechanical measurements on the stereocilia. Typically, the bundle of hairs moves as a unit, and the stiffness is in the range of $\kappa \sim 10^{-3}$ N/m or less. This implies that the Brownian motion of the bundle should have an amplitude $\delta x_{\text{rms}} = \sqrt{k_B T / \kappa} \sim 2$ nm. This seems small (remember that the stereocilia have lengths measured in microns), but ...




FIG. 76 Hair cells of the vertebrate inner ear [find better images, with scale bars!]. At left, in the bullfrog sacculus, from <http://www.hhmi.org/senses/c120.html>. At right, in the mammalian cochlea, three rows of “outer” hair cells and one row of “inner” hair cells at top, from Dallos (1984).

There is a particular species of neotropical frog, for example, that exhibits clear behavioral responses to vibrations of the ground that have an amplitude of ~ 1 Å. Individual neurons which carry signals from the hair cells in the sacculus the to brain actually *saturate* in response to vibrations of just ~ 10 Å = 1 nm. Although there are controversies about the precise numbers, the motions of our eardrum in response to sounds we can barely hear are similarly on the atomic scale. Invertebrates don’t use hair cells, but they also have mechanical sensors, and many of these too respond reliably to motions in the Ångström or even sub-Ångström range.

By itself, the order-of-magnitude (or more) discrepancy between the amplitude of Brownian motion and the threshold of sensation might or might not be a problem (we’ll come back to this). But surely it motivates us to ask if, by analogy with kinetic proofreading, it is possible to lower the effective noise level by pushing the system away from thermal equilibrium. This also is an interesting physics problem, independent of its connection to biology.

Consider a mass hanging from a spring, subject to drag as it moves through the surrounding fluid, as in Fig 77. By itself, the dynamics of this system are described by

the Langevin equation, [point back!]

$$m \frac{d^2x(t)}{dt^2} + \gamma \frac{dx(t)}{dt} + \kappa x(t) = F_{\text{ext}}(t) + \zeta(t), \quad (448)$$

where F_{ext} denotes external forces acting on the system and the Langevin force obeys

$$\langle \zeta(t) \zeta F(t') \rangle = 2\gamma k_B T \delta(t - t'). \quad (449)$$

But suppose that we measure the position of the mass, differentiate to obtain the velocity, and then apply a “feedback” force proportional to this velocity, $F_{\text{feedback}} = -\eta dx(t)/dt$; then we have

$$m \frac{d^2x(t)}{dt^2} + \gamma \frac{dx(t)}{dt} + \kappa x(t) = F_{\text{ext}}(t) + \zeta(t) + F_{\text{feedback}}(t) \quad (450)$$

$$= F_{\text{ext}}(t) + \zeta(t) - \eta \frac{dx(t)}{dt} \quad (451)$$

$$m \frac{d^2x(t)}{dt^2} + (\gamma + \eta) \frac{dx(t)}{dt} + \kappa x(t) = F_{\text{ext}}(t) + \zeta(t). \quad (452)$$

This system is equivalent to one with a new drag coefficient $\gamma' = \gamma + \eta$. But the fluctuating force hasn't changed—the molecules of the fluid don't know that we are applying feedback—so we can write

$$m \frac{d^2x(t)}{dt^2} + \gamma' \frac{dx(t)}{dt} + \kappa x(t) = F_{\text{ext}}(t) + \zeta(t) \quad (453)$$

$$\langle \zeta(t) \zeta F(t') \rangle = 2\gamma k_B T \delta(t - t') = 2\gamma' k_B T_{\text{eff}} \delta(t - t'), \quad (454)$$

where $T_{\text{eff}} = T\gamma/\gamma' = T\gamma/(\gamma + \eta)$. Thus, by observing the system and applying a feedback force, we synthesize a system which is, effectively, colder and thus has (in some obvious sense, but we will need to be careful) less thermal noise.




FIG. 77 A schematic of active feedback, in which we observe the position of a mass on a spring and apply a force proportional to the velocity. This can serve to enhance or compensate the intrinsic drag γ , but since it is generated by an active mechanism (symbolically, through the amplifier) there need not be an associated change in the magnitude of the Langevin force, as there would be at thermal equilibrium.

This idea of “active cooling” is very old, but it has received new attention in the attempt to build very sensitive displacement detectors, e.g. for the detection of gravitational waves. A recent example placed a one gram mass in a laser interferometer and used the change in radiation pressure on the mass as function of its position to generate the feedback force; this is different in detail from the model above, but similar in spirit. The result was that the effective temperature could be brought down from ~ 300 K to $\sim 7 \times 10^{-3}$ K, a reduction of roughly 40,000 \times , and this seems to be limited by noise in the laser itself.

It is important to be clear about exactly which measures of noise are reduced, and which are not. The mean-square displacement of the oscillator—and hence, by equipartition, the apparent temperature—has been reduced. But when we try to drive the system with a force at the resonant frequency, the added damping means that it is more resistant, and hence the response to a given force is smaller. Thus if we ask for the minimum force that we must apply (on resonance) to displace the oscillator by one standard deviation, this threshold force actually goes up, as if the system had more noise, not less. Finally, if we imagine that we can observe the position of the oscillator over a very long time, then what matters for detecting a small applied force at the resonant frequency is the spectral density of force noise, and this hasn't changed at all.

Problem 82: Effective noise levels. Do the real calculations required to verify the statements in the previous paragraph. These are not difficult.

As alternative to actively damping the oscillator, we can try to actively *undamp*, using feedback of opposite sign:

$$m \frac{d^2 x(t)}{dt^2} + \gamma \frac{dx(t)}{dt} + \kappa x(t) = F_{\text{ext}}(t) + \zeta(t) + F_{\text{feedback}}(t) \quad (455)$$

$$= F_{\text{ext}}(t) + \zeta(t) + \eta \frac{dx(t)}{dt} \quad (456)$$

$$m \frac{d^2 x(t)}{dt^2} + (\gamma - \eta) \frac{dx(t)}{dt} + \kappa x(t) = F_{\text{ext}}(t) + \zeta(t). \quad (457)$$

Now the variance of the displacement is larger,

$$\langle (\delta x)^2 \rangle = \frac{k_B T_{\text{eff}}}{\kappa} = \frac{k_B T}{\kappa} \cdot \frac{\gamma}{\gamma - \eta}, \quad (458)$$

but the sensitivity to forces applied on resonance is also enhanced. If we have $F_{\text{ext}}(t) = F_0 \cos(\omega_0 t)$, with $\omega_0 = \sqrt{\kappa/m}$, then the displacement will be $x(t) = x_0 \sin(\omega_0 t)$, with $x_0 = F_0 / [(\gamma - \eta)\omega_0]$. Thus the signal-to-noise ratio in a snapshot of the motion becomes

$$\frac{x_0^2}{\langle (\delta x)^2 \rangle} = \frac{F_0^2}{(\gamma - \eta)^2 \omega_0^2} \cdot \frac{\gamma - \eta}{\gamma} \frac{\kappa}{k_B T} = \left[\frac{\kappa F_0^2}{(\gamma \omega_0^2)^2 k_B T} \right] \cdot \frac{\gamma}{\gamma - \eta}. \quad (459)$$

Thus, in this case the signal-to-noise ratio for a snapshot of the position goes up in proportion to the amount of active ‘undamping.’

We can understand the impact of active undamping as a narrowing of the system bandwidth, or a sharpening of the resonance around ω_0 . Both the external force and the Langevin force drive the system in the same way. The difference is that we are considering an external force at the resonant frequency, while the Langevin force is white noise, with equal power at all frequencies. By sharpening the resonance, active undamping reduces the total impact of this noise; since the bandwidth of the resonance is proportional to $\gamma - \eta$, the enhancement of the signal-to-noise ratio is also in proportion to this factor.

Taken at face value it seems that we can increase the signal-to-noise ratio by an arbitrarily large factor—if we increase η so that $\gamma - \eta \rightarrow 0$, the resonance becomes infinitely sharp and it becomes possible to detect arbitrarily small forces from just an instantaneous look at the position x . Any recipe for detecting arbitrarily small signals should be suspect, but what actually limits the growth of the signal-to-noise ratio in this case?

First, it should be clear that the increased SNR comes at a cost. In a system with a sharp resonance, the time scale for response becomes long in inverse proportion to the bandwidth. Thus, as we let $\gamma - \eta \rightarrow 0$, the current position $x(t)$ becomes dependent on the forces $F_{\text{ext}}(t)$ in distant past. This is a serious issue, but it doesn’t really

set a limit to the smallest force we can detect.

Problem 83: A reminder about Green functions. The solution to the equation

$$m \frac{d^2 x(t)}{dt^2} + (\gamma - \eta) \frac{dx(t)}{dt} + \kappa x(t) = F_{\text{ext}}(t) \quad (460)$$

can be written in the form

$$x(t) = \int dt' G(t - t') F_{\text{ext}}(t'), \quad (461)$$

where $G(\tau)$ is the Green function or (time domain) linear response function. Find $G(\tau)$, and verify that as $\gamma - \eta \rightarrow 0$ this function acquires weight at very large τ , corresponding to a very long memory or strongly nonlocal responses.

A second limit to the signal-to-noise ratio is set by noise in the amplifier itself. This certainly is a practical problem, and there may even be a fundamental problem, since linear amplifiers have a minimum level of noise set by quantum mechanics. There is some very interesting physics here, and (confession time) there was a time when I worked very hard to convince myself that these quantum limits to measurement could be relevant to biological systems. This project failed, and I would rather not revisit old failures, so let’s skip this one.

The third consideration which limits the narrowing of the bandwidth is the finite power output of any real amplifier. As we let $\gamma - \eta \rightarrow 0$, the amplitude of motion in response to a force at resonance grows as $1/(\gamma - \eta)$, and since there is a real drag force $-\gamma(dx/dt)$ the amplifier must dissipate power to drive these ever larger motions. At some point this power requirement will become overwhelming, and the simple model $F_{\text{feedback}} = +\eta(dx/dt)$ has to break down. Intuitively, we expect that as x be-

comes larger, the strength of the feedback will decrease, so we can describe at least the beginning of this power limitation we can write

$$\eta \rightarrow \eta(x) \approx \eta_0[1 - (x/x_s)^2 + \dots], \quad (462)$$

where x_s is the scale on which the amplifier loses linearity. Then we have

$$m \frac{d^2x(t)}{dt^2} + (\gamma - \eta_0) \frac{dx(t)}{dt} + \frac{\eta_0}{x_s^2} x^2(t) \frac{dx(t)}{dt} + \kappa x(t) = F_{\text{ext}}(t) + \delta F(t). \quad (463)$$

This equation has several important features.

First, $\gamma = \eta_0$ is a bifurcation point. If $\gamma > \eta_0$, then in the absence of forces any small displacement from $x = 0$ will decay with time. In contrast, for $\gamma < \eta_0$, small displacements will oscillate and grow until the nonlinear term $\sim x^2(dx/dt)$ becomes significant. This is an example of a Hopf bifurcation [should we say some more technical things here about the kinds of bifurcations and the defining features of Hopf?]. Second, if we poise the system precisely at the bifurcation point, and drive it with a resonant force, then neglecting noise we have

$$m \frac{d^2x(t)}{dt^2} + \frac{\gamma}{x_s^2} x^2(t) \frac{dx(t)}{dt} + \kappa x(t) = F_0 \cos(\omega_0 t). \quad (464)$$

Guessing that the solution is of the form $x(t) \approx x_0 \sin(\omega_0 t)$, we note that

$$x^2(t) \frac{dx(t)}{dt} \approx \omega_0 x_0^3 \sin^2(\omega_0 t) \cos(\omega_0 t) \quad (465)$$

$$= \frac{1}{4} \omega_0 x_0^3 [\cos(\omega_0 t) - \cos(3\omega_0 t)]; \quad (466)$$

in the limit that the resonance is sharp, we know that the term at frequency $3\omega_0$ can't really drive the system, so we neglect this. Thus we have

$$\frac{\gamma \omega_0}{4x_s^2} x_0^3 = F_0, \quad (467)$$

or

$$x(t) = \left[\frac{4F_0 x_s^2}{\gamma \omega_0} \right]^{1/3} \sin(\omega_0 t). \quad (468)$$

Thus, the response to applied forces is nonanalytic (at least in the absence of noise); the slope of the response at $F_0 = 0$ is infinite, as one expects from the linear equation above, but the response to any finite force is finite.

The fractional power behavior in Eq (468) connects to a well known but very puzzling fact about the auditory system. As with any nonlinear system, if we stimulate

the ear with sine waves at frequencies f_1 and f_2 , we can hear “combination tones” built out of these fundamentals: $f_1 \pm f_2$, $2f_1 - f_2$, and so on. In the human ear, the term $2f_1 - f_2$ (with $f_1 < f_2$) is especially prominent. What is surprising is that the subjective intensity of this combination tone is proportional to the intensity of the fundamental tones. If we imagine that combination tones arise from a weak nonlinearity that could be treated in perturbation theory, we would predict that if the input tones have amplitudes A_1 and A_2 , then the amplitude of the combination tone should be $A_{2f_1 - f_2} \propto A_1^2 A_2$. In contrast, the model poised precisely at the bifurcation point predicts $A_{2f_1 - f_2} \propto (A_1^2 A_2)^{1/3}$, so that if we double the intensity of the input sounds we also double the intensity of the combination tone, as observed.

Problem 84: Combination tones. Do honest calculations to verify the statements about combination tones in the previous paragraph. Contrast the predictions far from the bifurcation point, where perturbation theory is applicable, with the predictions at the bifurcation point.

What happens to the nominally infinite signal-to-noise ratio in the linear model? As we increase the feedback η , the mean square displacement increases, but Eq (462) tells us that at larger x the effective strength of the feedback term decreases. We can try to see what will happen by asking for self-consistency. Suppose we replace the x -dependent value of the feedback term by an effective feedback strength which is given by the average,

$$\eta_{\text{eff}} \equiv \langle \eta(x) \rangle = \eta_0 [1 - \langle x^2 \rangle / x_s^2]. \quad (469)$$

But if we have an effective feedback term we can go back to the linear problem, and then Eq (458) tells us that

$$\langle x^2 \rangle = \frac{k_B T}{\kappa} \cdot \frac{\gamma}{\gamma - \eta_{\text{eff}}}. \quad (470)$$

Combining these equations gives us a self-consistent equation for the position variance $\langle x^2 \rangle$,

$$\frac{\eta_0}{\gamma x_s^2} \langle x^2 \rangle^2 + \left(1 - \frac{\eta_0}{\gamma}\right) \langle x^2 \rangle = \frac{k_B T}{\kappa}. \quad (471)$$

Even if we let the strength of the bare feedback η_0 become infinitely large, this equation predicts that the effective feedback term will remain finite, and in particular we always have $\eta_{\text{eff}} < \gamma$, so we can never cross the bifurcation, at least in this approximation. Concretely, solving Eq (471) and substituting back into Eq (469) for the effective feedback, we find

$$\lim_{\eta_0 \rightarrow \infty} \frac{\gamma - \eta_{\text{eff}}}{\gamma} = \frac{k_B T}{\kappa x_s^2}. \quad (472)$$

Thus, the system can narrow its bandwidth to an extent that is limited by the dynamic range of the feedback amplifier, which in turn is related to its power output. Since active narrowing of the bandwidth reduces the effective noise level below the expected thermal noise, we have a situation every much analogous to kinetic proofreading: we can do better than Boltzmann, but it costs energy, and the more energy the system expends, the better it can do.

Problem 85: Noise levels in nonlinear feedback. Start by verifying Eq (472). In the same approximation, calculate the response to applied forces, and show that the smallest force which be detected above the noise has been reduced by a factor $\sim \sqrt{\kappa x_s^2 / k_B T}$ relative to what we would have without feedback. Then, there are several things to worry about.

(a.) We have given two analyses. In the first, leading to Eq (468), we neglect noise and take the nonlinearities seriously, finding that the response to small forces is non-analytic. In the second, leading to Eq (472) we treat the crucial nonlinear terms as a self-consistently determined linear feedback, and noise is included. In this second approach, the response to applied forces is linear. Can you reconcile these approaches? Presumably the first approach is valid if the applied forces produce displacements much larger than the noise level. Does this mean that the noise serves to “round” the nonanalytic behavior near $F = 0$?

(b.) How do your results in (a.) effect your estimates of the smallest force that can be detected above the noise?

(c.) You might be worried that our self-consistent approximation is a bit crude. An alternative is to simulate Eq (463) numerically, reminding yourself of the discussion in Section II.A about how to treat the Langevin force. Compare the results of your simulation with the predictions of the self-consistent approximation, for example Eq (471).

(d.) You could also try an alternative analytic approach. If we rewrite Eq (463) in the absence of external forces as

$$\begin{aligned} \frac{dx(t)}{dt} &= v(t) & (473) \\ m \frac{dv(t)}{dt} &= - \left[\gamma - \eta_0 \left(1 - \frac{x^2(t)}{x_s^2} \right) \right] v(t) - \kappa x(t) + \delta F(t), & (474) \end{aligned}$$

then you should be able to derive a Fokker-Planck or diffusion-like equation for the probability $P(x, v)$ of finding the system with

instantaneous position x and velocity v . Can you find the steady state solution? How does this compare with your numerical results?

What do we learn from all this? Although there are limits, active feedback (with either sign) makes it possible to detect smaller signals than might otherwise be possible given the level of thermal noise. Pushing the system away from equilibrium, we spend energy to improve performance. This sounds like the sort of thing biological systems might exploit.

If thermal noise is important, then it is useful to think about the bandwidth the system is using as it “listens” (in this case, literally) to its input, and the resulting exchange of energy. We recall that in a resonator, the time scale on which oscillations decay away is $\tau \sim 1/\Delta f$, where Δf is the range of frequencies under the resonant peak. Thus if we excite the resonator to an amplitude such that it stores energy E , this energy also decays away on a time scale $\sim \tau$. But in thermal equilibrium we know that the average energy is not zero, but rather $k_B T$, so the surrounding heat bath must provide a flux of power $\sim k_B T / \tau \sim k_B T \Delta f$ to balance the dissipation. If we want to detect incoming signals above the background of thermal noise, then these signals have to deliver a comparable amount of power. A more careful calculation shows that this “thermal noise power” is $P = 4k_B T \Delta f$.

Problem 86: Acoustic cross-sections and detailed balance. Use idea of thermal noise power to derive limit on absorption cross-section averaged over directions. Emphasize connection to Einstein’s argument about A and B coefficients. Maybe look at data on the ear in relation to this limit?

Estimates of the power entering the inner ear at the threshold of hearing are $P \sim 4 \times 10^{-19}$ W. This suggests that, to be sure the signals are above thermal noise, the ear must operate with a bandwidth of less than $\Delta f \sim 100$ Hz. There are several ways of seeing that this is about right. If we record the responses of individual neurons emerging from the cochlea of animals like us, and we can see that these responses are tuned. More quantitatively, as in Fig [**], we can measure the sound pressure required to keep the neuron generating spikes at some fixed rate, and see how this varies with the frequency of pure tone inputs. This input required for constant output is minimal at one “characteristic frequency” of the neuron, and rises steeply away from this minimum; for neurons with characteristic frequencies in the range of 1 kHz, the bandwidths are indeed $\Delta f \sim 100$ Hz. One can

also try to measure the effective bandwidth in human observers, either by asking listeners to detect a tone in a noisy background and seeing how detection performance varies with the width of the noise, or by testing when one tone impairs the detection of another. More recently it has been possible to record the responses from individual receptor cells, as in Fig [This paragraph needs figures with recordings from primary auditory neurons and hair cells; be sure that these are properly referenced at the end of the section.] All of these bandwidth estimates are in rough agreement, and also agree with the estimate based on comparing thermal noise with the power entering the ear at threshold, suggesting that filtering—in addition to its role in decomposing sounds into their constituent tones—really is essential in limiting the impact of noise. It is important that the resonance or filter which defines this bandwidth actually be in a part of the system where it can act to reject the dominant source of thermal noise. For example, if we think of the vibration sensitive frog, placing the frog on a resonant table would mean that the whole system had a narrower bandwidth, but this would do nothing to reduce the impact of random motions of the stereocilia. It is extremely implausible that the passive mechanics of the stereocilia themselves can generate this narrow bandwidth.

Problem 87: Stereocilium mechanics. Use the image of the hair bundle in Fig ** to estimate the mass and drag coefficient of the bundle as it moves through the surrounding fluid, which you can assume is water. Is the system naturally resonant? Overdamped or underdamped? What bandwidth of filtering would be needed to be sure that fluid displacements of $\sim 1 \text{ \AA}$ are detectable above the thermal noise of the bundle? Is this roughly consistent with the observed threshold power?

In mammalian ears, the hair cells sit on top of a structure called the basilar membrane, the tips of the stereocilia are in contact with another structure, the tectorial membrane, and the entire organ, called the cochlea, is wrapped into a spiral and embedded in bone [need a figure here!]. Sound waves impinging on the eardrum are coupled into the cochlea to produce a pressure difference across the basilar membrane, which then vibrates, ultimately causing motions of the stereocilia. Because it is surrounded by fluid, motions of neighboring pieces of the basilar membrane are coupled, and the result is a wave that travels along the membrane; because of gradations in the mechanical properties of the system, high frequency waves have their peak amplitude near the entrance to the cochlea and low frequency waves have their peak near the end or apex of the cochlea. Helmholtz knew about the structure of the inner ear, and since he

saw fibrous components in the various membranes, he imagined that these might be taught, resonant strings. Because the strings were of different lengths and thicknesses, varying smoothly along the length of the cochlea, the resonant frequency would also vary. Thus, Helmholtz had the basic picture of the cochlea as a mechanical system which analyzes incoming sounds into component frequencies, sorting them to different locations along the basilar membrane. It is not clear how seriously he took the details of the mechanics, but the picture of the ear as frequency analyzer or bank of filters was taken very seriously, and indeed this picture accounts for many perceptual phenomena. The first direct measurements of basilar membrane motion were made by von Békésy, who opened the cochleae of various animals, sprinkled reflecting flakes onto the membrane, and observed its motion stroboscopically under the microscope.⁵⁷ Békésy saw the traveling wave of vibrations along the basilar membrane, and he saw the mechanical sorting of frequencies which Helmholtz had predicted.

Problem 88: Cochlear mechanics. Generate a problem that gives the students a tour of classical ideas about the traveling wave along the basilar membrane. Get them to use WKB methods to solve, understand how the peak forms etc..

Békésy was also immediately impressed with the scale of motions in the inner ear. To make the basilar membrane vibrate by $\sim 1 \mu\text{m}$ and hence be easily visible under the light microscope, he had to deliver sounds at what would be the threshold of pain, $\sim 120 \text{ dB SPL}$.⁵⁸ If we just extrapolate linearly, $1 \mu\text{m}$ at 120 dB SPL corresponds to 10^{-12} m at 0 dB SPL , or $\sim 0.01 \text{ \AA}$ (!). This is an astonishingly small displacement.

⁵⁷ Many of von Békésy's key contributions are collected in a volume published relatively late in his life, along with various reminiscences and quasi-philosophical remarks. As an example, he notes that in science good enemies are much more valuable than good friends, since enemies will take the time to find all your mistakes. Unfortunately, in the process of this dialogue, some of the enemies become friends and hence, by von Békésy's criteria, their usefulness is lost.

⁵⁸ SPL stands for sound pressure level. It is conventional in acoustics to measure the intensity of sounds logarithmically relative to some standard. 10 dB corresponds to a power ratio of $10\times$, so 20 dB corresponds to a factor of $10\times$ higher sound pressure variations. For human hearing the standard reference (0 dB SPL) is a pressure of $2 \times 10^{-5} \text{ N/m}^2$ which is close to the threshold of hearing at frequencies near 2 kHz .

Problem 89: Brownian motion of the basilar membrane.

Generate a problem that takes the students through the analysis of Brownian motion in a continuous system, with basilar membrane as an example.

Békésy also observed that the frequency selectivity of the basilar membrane motion was quite modest. More precisely, the peak of the vibrations in response to a single frequency was quite broad, spreading over a distance along the cochlea that corresponds to more than ten times the apparent bandwidth over which we integrate. This discrepancy seems to have caused more concern than the extrapolated displacement. On the one hand, if it is correct it suggests that there are mechanisms to sharpen frequency selectivity that come after the mechanics of the inner ear, perhaps at the level of neural circuitry. Békésy was very much taken with the ideas of lateral inhibition in the retina, and suggested that this might be a much more general concept for neural signal processing. On the other hand, von Békésy studied dissected cochleae that were, not to put too fine a point on it, dead. By the 1970s, it became clear that individual neurons emerging from the cochlea had frequency selectivity which was sharper than suggested by von Békésy's measurements, and that (especially in mammals) this selectivity was extremely fragile, dependent on the health of the cochlea—so much so that the tuning properties of individual neurons could be changed within minutes by blocking blood flow to the ear, recovering just as quickly when the block was relieved.

Observations on the fragility of cochlear tuning emphasized the challenge of making direct mechanical measurements on more intact preparations, and presumably at more comfortable sound levels. To make measurements of smaller displacements, a number of tools from experimental physics were brought to bear: the Mössbauer effect, laser interferometry, and Doppler velocimetry. At the same time, several groups turned to non-mammalian systems which seemed like they would be more robust, such as the frog sacculus and the turtle cochlea, and especially in these systems it proved possible to make much more quantitative measurements on the electrical responses of the hair cells and eventually on their mechanical properties. In the midst of all this progress came the most astonishing evidence for active mechanical filtering in the inner ear.

If we build an active filter via feedback, and try to narrow the bandwidth as much as possible, we are pushing the system to the edge of instability. It is not difficult to imagine that, with active feedback provided by biological mechanisms, that some sort of pathology could result in an error that pushed past the gain past the bifurcation, turning a narrow bandwidth filter into an oscillator. If incoming sounds are efficiently coupled to motions of




FIG. 78 Spontaneous emission of sounds from the human ear, from van Dijk et al (2011). Top panels show the spectral density of sounds in the ear canals of two subjects. Bottom panel shows the intensities and frequencies of 41 spectral peaks found in 8 subjects, compared with the noise background.

the active elements in the inner ear, then spontaneous oscillations of these elements will couple back, and the ear will emit sound. Strange as it may seem, careful surveys show that almost half of all ears have a “spontaneous oto-acoustic emission;” a rather quiet, narrow band sound that can be detected by placing a microphone in the ear canal, as shown in Fig 78. Importantly, the statistics of the sounds being emitted are not those of filtered noise, but rather those expected from a true oscillator—the distribution of instantaneous sound pressures has a minimum at zero, as expected if the quiet state is unstable.

[Need to wrap this up .. Direct measurement on ciliary mechanics in different systems; violation of FDT as evidence of activity. Note re electrical resonances. Look at Marcelo & Jim's papers to see smoking gun for Hopf bifurcation.]

[Reach a conclusion!]

Now is a good time to look back at Schrödinger's remarkable little book (Schrödinger 1944). The idea which him were presented by Timoféef-Ressovsky et al (1935). For some later perspectives see Delbrück's Nobel lecture (1970); the title refers to an earlier lecture, also very much worth reading for its eloquence and prescience (Delbrück 1949). A review of DNA structure is given in Appendix A.5, and some general references on molecular biology are at the end of Section II.B. The ideas of kinetic proofreading—and, as emphasized in the text, the idea that there is a general physics problem cutting across a wide range of biological phenomena—were presented in Hopfield (1974) and Ninio (1975). Hopfield (1980) constructed a scenario in which the basic idea of paying (energetically) for increased accuracy still operates, but with none of the experimental signatures of the original proofreading scheme. [Need refs that proofreading is correct!]

- Delbrück 1949:** A physicist looks at biology. M Delbrück, *Trans Conn Acad Arts Sci* **38**, 173–190 (1949). Reprinted in *Phage and the Origins of Molecular Biology*, J Cairns, GS Stent & JD Watson, eds, pp 9–22 (Cold Spring Harbor Press, Cold Spring Harbor NY, 1966).
- Delbück 1970:** A physicist's renewed look at biology: twenty years later. M Delbrück, *Science* **168**, 1312–1315 (1970). Also available at <http://nobelprize.org>.
- Hopfield 1974:** Kinetic proofreading: A new mechanism for reducing errors in biosynthetic processes requiring high specificity. JJ Hopfield, *Proc Nat'l Acad Sci (USA)* **71**, 4135–4139 (1974).
- Hopfield 1980:** The energy relay: a proofreading scheme based on dynamic cooperativity and lacking all characteristic symptoms of kinetic proofreading in DNA replication and protein synthesis. JJ Hopfield, *Proc Nat'l Acad Sci (USA)* **77**, 5348–5252 (1980).
- Ninio 1975:** Kinetic amplification of enzyme discrimination. J Ninio *Biochimie* **57**, 587–595 (1975).
- Schrödinger 1944:** *What is Life?* E Schrödinger (Cambridge University Press, Cambridge, 1944).
- Timoféef-Ressovsky et al 1935:** Über die Natur der Genmutation und der Genstruktur. NW Timoféef-Ressovsky, KG Zimmer & M Delbrück, *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen* **1**, 190–245 (1935). [\[Is there a translation available?\]](#)
- Need refs for mutator strains, selection on error rates, tradeoff among energy costs, accuracy and speed of growth, ...** The fact that some mutations lead to changes in mutation rate could have dramatic consequences for the pace of evolutionary change, as emphasized by Magnasco & Thaler (1996). [\[where they "right," even in outline? what has happened since?\]](#)
- Magnasco & Thaler 1996:** Changing the pace of evolution. MO Magnasco & DS Thaler, *Phys Lett A* **221**, 287–292 (1996).
- The basic idea of kinetic proofreading—an enzymatic mechanism dissipating energy to stabilize a “better than Boltzmann” distribution of molecular states—has by now been applied in several different contexts. For the disentangling of DNA strands, see (Yan et al 1999, 2001), who were inspired in part by the experiments of Rybenkov et al (1997). For the sensitivity and specificity of initial events in the immune response, see McKeithan (1995) and Altan-Bonnet & Germain (2005), and for a more general view of signal transduction specificity see Swain & Siggia (2002).
- Altan-Bonnet & German 2005:** Modeling T cell antigen discrimination based on feedback control of digital ERK responses. G Altan-Bonnet & RN Germain, *PLoS Biology* **3**, 1925–1938 (2005).
- McKeithan 1995:** Kinetic proofreading in T-cell receptor signal transduction. TW McKeithan, *Proc Nat'l Acad Sci (USA)* **92**, 5042–5046 (1995).
- Rybenkov et al 1997:** Simplification of DNA topology below equilibrium values by type II topoisomerases. VV Rybenkov, C Ullsperger, AV Vologodskii & NR Cozzarelli, *Science* **277**, 690–693 (1997).
- Swain & Siggia 2002:** The role of proofreading in signal transduction specificity. PS Swain & ED Siggia, *Biophys J* **82**, 2928–2933 (2002).
- Yan et al 1999:** A kinetic proofreading mechanism for disentanglement of DNA by topoisomerases. J Yan, MO Magnasco & JF Marko, *Nature* **401**, 932–935 (1999).
- Yan et al 2001:** Kinetic proofreading can explain the suppression of supercoiling of circular DNAs by type-II topoisomerases. J Yan, MO Magnasco & JF Marko, *Phys Rev E* **63**, 031909 (2001).
- The use of optical forces to manipulate biological systems goes back to work by Ashkin & Dziedzic (1987), which in turn grew out of Ashkin's earlier work (Ashkin 1978, 1980). It is worth noting that the same ideas of optical trapping for neutral, dielectric particles were a key step in the development of atomic cooling, as described, for example, by Chu (2002). For the state of the art in single molecule experiments, see Greenleaf et al (2007). The experiments on single RNA polymerase molecules were by Shaevitz et al (2003) and Abbondanzieri et al (2005). [\[Add refs for single ribosome experiments, depends on what happens in the text.\]](#)
- Abbondanzieri et al 2005:** Direct observation of base-pair stepping by RNA polymerase. EA Abbondanzieri, WJ Greenleaf, JW Shaevitz, R Landick & SM Block, *Nature* **438**, 460–465 (2005).
- Ashkin 1978:** Trapping of atoms by resonance radiation pressure. A Ashkin, *Phys Rev Lett* **40**, 729–732 (1978).
- Ashkin 1980:** Applications of laser radiation pressure. A Ashkin, *Science* **210**, 1081–1088 (1980).
- Ashkin & Dziedzic 1987:** Optical trapping and manipulation of viruses and bacteria. A Ashkin & JM Dziedzic, *Science* **235**, 1517–1520 (1987).
- Chu 2002:** The manipulation of neutral particles. S Chu, in *Nobel Lectures, Physics 1996–2000* G Ekspung, ed (World Scientific, Singapore, 2002). Also available at <http://nobelprize.org>.
- Greenleaf et al 2007:** High-resolution, single-molecule measurements of biomolecular motion. WJ Greenleaf, MT Woodside & SM Block. *Annu Rev Biophys Biomol Struct* **36**, 171–190 (2007).
- Shaevitz et al 2003:** Backtracking by single RNA polymerase molecules observed at near-base-pair resolution. JW Shaevitz, EA Abbondanzieri, R Landick & SM Block, *Nature* **426**, 684–687 (2003).
- The images of hair cells are from [\[be careful to revise with new figures\]](#) and Dallos (1984). Early experiments on the mechanics of the stereocilia are by Flock & Strelhoff (1984; Strelhoff & Flock 1984). The remarkable vibration sensitivity of the frog inner ear is described in two papers, Narins & Lewis (1984) and Lewis & Narins (1985). Estimates of the power flowing into the inner ear at threshold, as well as the minimum detectable power in other sensory systems, were given some time ago by Khanna & Sherrick (1981). [\[I think there is also a nice reference to threshold powers in relation to electroreceptors; try to find this.\]](#)
- Dallos 1984:** Peripheral mechanisms of hearing. P Dallos, in *Comprehensive Physiology 2011, Supplement 3: Handbook of Physiology, The Nervous System, Sensory Processes* pp. 595–637 (Wiley-Blackwell, 1984).
- Flock & Strelhoff 1984:** Studies on hair cells in isolated coils from the guinea pig cochlea. Å Flock & D Strelhoff, *Hearing Res* **15**, 11–18 (1984).
- Khanna & Sherrick 1981:** The comparative sensitivity of selected receptor systems. SM Khanna & CE Sherrick, in *The Vestibular System: Function and Morphology*, T Gualtierotti, ed, pp 337–348 (Springer-Verlag, New York, 1981).
- Lewis & Narins 1985:** Do frogs communicate with seismic signals? ER Lewis & PM Narins, *Science* **227**, 187–189 (1985).
- Narins & Lewis 1984:** The vertebrate ear as an exquisite seismic sensor. PM Narins & ER Lewis, *J Acoust Soc Am* **76**, 1384–1387 (1984).
- Strelhoff & Flock 1984:** Stiffness of sensory-cell hair bundles in the isolated guinea pig cochlea. D Strelhoff & Å Flock, *Hearing Res* **15**, 19–28 (1984).

[Find early references to active cooling ...] Recent examples of active cooling are by Corbitt et al (2007) and Abbott et al (2009), who are aiming at improving the sensitivity of gravitational wave detection. For discussion of the quantum limits to mechanical measurements see Caves et al (1980), Caves (1985), and Braginsky & Khalili (1992). For the quantum limits to amplifier noise (which has a long history), see Caves (1982). For discussions of thermal noise, I have always liked the treatment in Kittel’s little book, cited at the end of Section 2.1, and this includes a discussion of the “thermal noise power.”

Abbott et al 2009: Observation of a kilogram-scale oscillator near its quantum ground state. B Abbott et al (LIGO collaboration), *New J Phys* **11**, 073032 (2009).

Braginsky & Khalili 1992: *Quantum Measurement*. VB Braginsky & F Ya Khalili (Cambridge University Press, Cambridge, 1992).

Caves 1982: Quantum limits on noise in linear amplifiers. CM Caves, *Phys Rev D* **26**, 1817–1839 (1982).

Caves 1985: Defense of the standard quantum limit for free-mass position. CM Caves, *Phys Rev Lett* **54**, 2465–2468 (1985).

Caves et al 1980: On the measurement of a weak classical force coupled to a quantum-mechanical oscillator. I: Issues of principle. CM Caves, KS Thorne, RWP Drever, VD Sandberg & M Zimmermann, *Revs Mod Phys* **52**, 341–392 (1980).

Corbitt et al 2007: Optical dilution and feedback cooling of a gram-scale oscillator to 6.9 mK. T Corbitt, C Wipf, T Bodiya, D Ottaway, D Sigg, N Smith, S Whitcomb & N Mavalvala, *Phys Rev Lett* **99**, 160801 (2007).

The physics of hearing is a subject that goes back to Helmholtz (1863) and Rayleigh (1877). An early measurement using tone-on-tone masking to probe the sharpness of frequency selectivity in the cochlea was done by Wegel & Lane (1924); a modern account of psychophysical measurements on effective detection bandwidths is given by [find ref!]. The classical measurements on cochlear mechanics by von Békésy were collected in 1960. The modern era of such measurements begins with the use of the Mössbauer effect (Frauenfelder 1962) to measure much smaller displacements of the basilar membrane (Johnstone et al 1967), and to demonstrate these motions have decided sharper frequency tuning in response to quieter sounds (Rhode 1971). These experiments triggered renewed interest in theories of cochlear mechanics, and some beautiful papers from this period are by Zweig et al (1976; Zweig 1976) and by Lighthill (1981), both of which connect the dynamics of the inner ear to more general physical considerations. There then followed a second wave of mechanical measurements, using interferometry (Khanna & Leonard 1982), Doppler velocimetry [Ruggero?], and improved Mössbauer methods (Sellick et al 1982). Some perspective on these measurements and models as they emerged can be found in Lewis et al (1985), which also makes connections to the mechanics of other inner ear organs. [Something more modern?]

Békésy 1960: *Experiments in Hearing* G von Békésy, EG Wever, ed (McGraw-Hill, New York, 1960).

Frauenfelder 1962: *The Mössbauer Effect* (WA Benjamin, New York, 1962).

Helmholtz 1863: *Die Lehre den Tonempfindungen als physiologische Grundlage für die Theorie der Musik*. H von Helmholtz (Vieweg and Sohn, Braunschweig, 1863). The most widely used translation is the second English edition, based on the fourth (and last) German edition of 1977; translated by AJ Ellis with an introduction by H Margenau, *On the Sensations of Tone as a Physiological Basis for the Theory of Music* (Dover, New York, 1954).

Johnstone et al 1967: Basilar membrane vibration examined with the Mössbauer technique. BM Johnstone & AJF Boyle, *Science* **158**, 390–391 (1967).

Khanna & Leonard 1982: Basilar membrane tuning in the cat cochlea. SM Khanna & DGB Leonard, *Science* **215**, 305–306 (1982).

Lewis et al 1985: *The Vertebrate Inner Ear* ER Lewis, EL Levrenz & WS Bialek (CRC Press, Boca Raton FL, 1985).

Lighthill 1981: Energy flow in the cochlea. J Lighthill, *J Fluid Mech* **106**, 149–213 (1981).

Rayleigh 1877: *Theory of Sound* JW Strutt, Baron Rayleigh (Macmillan, London, 1877). More commonly available is the revised and enlarged second edition, with a historical introduction by RB Lindsay (Dover, New York, 1945).

Rhode 1971: Observations of the vibration of the basilar membrane in squirrel monkeys using the Mössbauer technique. WS Rhode, *J Acoust Soc Am* **49**, 1218–1231 (1971).

Sellick et al 1980: Measurement of basilar membrane motion in the guinea pig using the Mössbauer technique. PM Sellick, R Patuzzi & BM Johnstone, *J Acoust Soc Am* **72**, 131–141 (1982).

Wegel & Lane 1924: The auditory masking of one pure tone by another and its probable relation to the dynamics of the inner ear. RL Wegel & CE Lane, *Phys Rev* **23**, 266–285 (1924).⁵⁹

Zweig 1976: Basilar membrane motion. G Zweig, *Cold Spring Harb Symp Quant Biol* **40**, 619–633 (1976).

Zweig et al 1976: The cochlear compromise. G Zweig, R Lipps & JR Pierce, *J Acoust Soc Am* **59**, 975–982 (1976).

The idea of active filtering in the inner ear goes back to a remarkably prescient paper by Gold (1948), who is better known, perhaps, for his contributions to astronomy and astrophysics; see Burbidge & Burbidge (2006). The idea that active elements are at work in the mechanics of the mammalian cochlea gained currency as experiments showed the “vulnerability” of frequency selectivity (Evans 1972), and with the dramatic observation of acoustic emissions from the ear (Kemp 1978, Zurek 1981); the data shown in Fig 78 are from van Dijk et al (2011). Importantly these emissions are observed not only from the rather complex mammalian cochlea, but from simpler ears of amphibians [get proper refs from van Dijk et al]. [Need a recent overview of acoustic emissions.] The idea that active filtering is essential for noise reduction is discussed in Bialek (1987). The modern view of active filtering as an approach to the Hopf bifurcation begins with Eguíliz et al (2000), and has been developed by [...].

Bialek 1987: Physical limits to sensation and perception. W Bialek, *Annu Rev Biophys Biophys Chem* **16**, 455–478 (1987).

Burbidge & Burbidge 2006: Thomas Gold, 1920–2004. G Burbidge & M Burbidge, *Bio Mem Nat’l Acad Sci (USA)* **88**, 1–15 (2006).

van Dijk et al 2011: The effect of static ear canal pressure on human spontaneous otoacoustic emissions: Spectral width as a measure of the intra-cochlear oscillation amplitude. P van Dijk, B Maat & E de Kleine, *J Assoc Res Otolaryngol* **12**, 13–28 (2011).

Eguíliz et al 2000: Essential nonlinearities in hearing. VM Eguíliz, M Ospeck, Y Choe, AJ Hudspeth & MO Magnasco, *Phys Rev Lett* **84**, 5232–5235 (2000).

⁵⁹ It is amusing to note that this paper is sometimes cited in the biological literature as having been published in the journal *Physiological Reviews*. Presumably this reflects authors or editors copying the reference to *Phys Rev* and “correcting” it to *Physiol Rev* without checking.

- Evans 1972:** The frequency response and other properties of single fibres in the guinea-pig cochlear nerve. EF Evans, *J Physiol (Lond)* **226**, 263–287 (1972).
- Gold 1948:** Hearing. II: The physical basis of the action of the cochlea. T. Gold, *Proc R Soc Lond Ser B* **135**, 492–498 (1948).
- Kemp 1978:** Stimulated acoustic emissions from within the human auditory system. DT Kemp, *J Acoust Soc Am* **64**, 1386–1391 (1978).
- Zurek 1981:** Spontaneous narrowband acoustic signals emitted by human ears. PM Zurek, *J Acoust Soc Am* **69**, 514–523 (1981).

E. Perspectives

Many of life’s phenomena exhibit a startling degree of reliability and precision. Organisms reproduce and develop with surprising predictability, and our own perceptual experience of the world feels certain and solid. On the other hand, when we look inside a single cell, or even at the activity of single neurons in brain, things look very noisy. Are the building blocks of biological behavior really so noisy? If so, how can we understand the emergence of reliability and certainty from all this mess?

Many of the problems faced by living organisms can be phrased as sensing, processing and responding to signals. If we look at a part of a system involved in such sensory tasks, we have to be careful in assessing noise levels. As a simple example, if we build a system in the lab that measures a small signal, and somewhere in this system there is an amplifier with very high gain, then surely we will find places in the circuitry where the voltage fluctuations are very large. Alternatively, there might be no gain, just a lot of noise. Thus, the variance of the noise at one point in the system, by itself, tells us nothing about its true degree of noisiness.

When we build sensors in the lab, we measure their noise performance by referring the noise to the input—estimating the noise level that would have to be added to the signals that we are trying to sense so as to account for the noise that we see at the output. This effective noise level is also the noise that limits the detectability of small signals, or the discriminability of signals that are very similar to one another. Importantly, for many sensors there are physical limits on this effective noise at the input, which allows us to put the noise performance on an absolute scale.

What we have done in this Chapter is to look at several instances in which it has been possible to carry out the program of “referring noise to the input” for increasingly complex biological systems. This is by far not a closed subject, and it is a minority of systems that have been analyzed in this way. Nonetheless, it is striking that, in so many disparate instances, the noise performance of biological systems indeed is close to the relevant physical limits. This of course is in the spirit of what we learned

from the case of photon counting in vision, but it seems much more general.

[I need to give some exegesis of this, and what it implies. Perhaps because I have spent so much time on these issues myself, I am having difficulty at the moment generating enough distance to be clear and objective (and not just to repeat what was said at the end of the previous chapter). So, I will need to come back to this. Sorry to leave things hanging in an important spot!]

III. NO FINE TUNING

Imagine making a model of all the chemical reactions that occur inside a cell. Surely this model will have many thousands of variables, described thousands of differential equations. If we write down this many differential equations with the right general form but choose the parameters at random, presumably the resulting dynamics will be chaotic. Although there are periodic spurts of interest in the possibility of chaos in biological systems, it seems clear that this sort of “generic” behavior of large dynamical systems is not what characterizes life. On the other hand, it is not acceptable to claim that everything works because every parameter has been set to just the right value—in particular these parameters depend on details that might not be under the cell’s control, such as the temperature or concentration of nutrients in the environment. More specifically, the dynamics of a cell depend on how many copies of each protein the cell makes, and one either has to believe that everything works no matter how many copies are made (within reason), or that the cell has ways of exerting precise control over this number; either answer would be interesting. This problem—the balance between robustness and fine tuning—arises at many different levels of biological organization. Our goal in this chapter is to look at several examples, from single molecules to brains, hoping to see the common themes. **[This seems to be the thinnest, and least well worked out of all the four main chapters. All advice is welcome!]**

Physics, especially theoretical physics, is the search for concise mathematical descriptions of Nature, and to a remarkable extent this search has been successful. The dirty laundry of this enterprise is that our mathematical descriptions of the world have parameters. In a sense, one mathematical structure describes several possible worlds, which would be somewhat different if the parameters were chosen differently. Sometimes this variety is a good thing—in condensed matter physics, for example, the different parameter values might correspond to genuinely different materials, all of which are experimentally realizable. On the other hand, if the predictions of the model are too sensitive to the exact values of the parameters, there is something vaguely unsatisfying about our claim to have explained things. Such strongly parameter-dependent explanations are often called “finely tuned,” and we have grown to be suspicious of fine tuning. Experience suggests that if parameters need to be set to precise (or somehow unnatural) values, then we are missing something in our mathematical description of Nature.⁶⁰

One needs, of course, to be cautious in identifying examples of fine tuning. As an example, many of the beautiful phenomena associated with solar eclipses depend on the fact that, seen from our vantage point on the earth, the angular size of the moon is almost exactly equal to the angular size of the sun. As far as we know, this is a coincidence, and isn’t connected to anything else. Presumably this coincidence (which, at certain times of year, occurs with $\sim 1\%$ precision) is related to the fact that there are many planets with moons—even more if we count the planets orbiting other stars—and we happen to live on one of them. Thus, we are sampling one out of many possibilities, and so rare things will happen. Similarly, elections sometimes turn on a surprisingly small number of votes, a tiny fraction of the total. This might seem like some sort of fine tuning,⁶¹ but it is also true that *most* elections do not have outcomes anywhere near the point of perfect balance among the outcomes. This is more obviously one of those cases in which we are sampling many examples, and finely tuned outcomes will happen, sometimes, by chance alone. What we need to worry about are cases in which fine tuning seems essential to make things work (unlike the moon/sun example), and where we see this in representative examples, or in all examples (unlike the elections). We’ll see plenty of these problematic cases.

In biological systems, there may be different reasons to be suspicious of fine tuning. On the one hand, for many processes what we call parameters are certainly dynamical variables on longer time scales (such as the number of copies of a protein), and there is widespread doubt that cells can regulate these dynamics precisely. More fundamentally, the parameters of biological systems are encoded in the genome, and in order for evolution to occur it seems necessary that, near to the genomes we see today, there must be genomes (and hence parameter values) which also generate functional organisms of reasonable fitness. These ideas have entered the literature as the need for robustness and evolvability. Note that while the physicist’s suspicion of fine tuning is a statement about the kind of explanation that we find satisfying, any attempt to enshrine robustness and evolvability as specifically biological principles involves hypotheses, either about the ability of cells to control their internal states or about the dynamics of evolution.

In this section we will look at several examples of the fine problem, starting at the level of single molecules and then moving “up” to the dynamics of single neurons, the internal states of single cells more generally, and networks

⁶⁰ At this point I usually try to remind the students of examples—the apparent vanishing of CP violation for the strong interaction, and the prediction of the axion as a solution to this problem, is a favorite. The cosmological constant is another one. Whether these remarks help depends on what the students have learned

in other courses. **Would it be good to make this explicit here? In the text or a footnote?**

⁶¹ We’ll leave aside, for this discussion, the disturbing possibility that vote totals *are* being tuned by some process that is separate from the actions of the voters themselves.




FIG. 79 The basic structure of amino acids and the peptide bond. At top, two amino acids. Different amino acids are distinguished by different groups \mathbf{R} attached to the α -carbon. Proteins are polymers of amino acids, and the chemical step in polymerization is the formation of the “peptide bond” by removal of a water molecule.

of neurons. As noted at the outset, these different biological systems are the subjects of non-overlapping literatures, and so part of what I hope to accomplish in this Chapter is to highlight the commonality of the physics questions that have been raised in these very different biological contexts.

A. Sequence ensembles

The qualitative ideas about robustness vs fine tuning can be made much more concrete by focusing on single protein molecules. We recall that proteins are heteropolymers of amino acids (Fig 79), each monomer along the polymer chain chosen from twenty possible amino acids (Fig 80). When we look at the proteins made by one particular organism, of course each protein has some particular sequence. If a typical protein is 200 amino acids long, then there are $(20)^{200} \sim 10^{260}$ possible sequences, out of which a bacterium might choose a few thousand,

and we choose a few tens of thousands. While different organisms do make slightly different choices, even if we sum over all life forms on earth we will find that real proteins occupy a very small fraction of the available volume in sequence space.

Proteins with different sequences fold up into different structures and carry out different functions. Thus, the sequence obviously matters. Yet, it can't be that the *exact* sequence matters, and this can be checked experimentally. Although some changes are disastrous (e.g., trying to bury a charged amino acid deep in the interior of the protein), many amino acid substitutions leave the structure and function of a protein almost completely unchanged, and many more generate quantitative modulations of function which could be useful in different environments or for closely related organisms. [Should add some figures with protein structures. Need pointer to Appendix A.5 discussing methods of structure determination. Also need to point out that the possible folds seem to be limited, which is another indication that not all details matter.]

Although protein function is tolerant to a wide range of sequence changes, not all sequences really make functional proteins. We can make this statement both as a theoretical result and as an experimental fact. Experimentally, we can synthesize proteins by choosing amino acids at random, and almost none of these will fold. As we will see below, we can even bias our choices at each site, trying to emulate a known family of proteins, and it still is true that if we choose each amino acid independently, most proteins don't fold.

As a crude theoretical model of a protein, we can coarse grain to keep track of the positions \mathbf{r}_i of each α -carbon atom (see Fig 79) along the chain, not worrying about the detailed configuration of the side chains that project from the backbone. Successive amino acids are bonded to one another, with a relatively fixed bond length ℓ , and when the chain folds to bring two amino acids near one another they have an interaction that depends on their identity, plus an excluded volume interaction that is independent of identity. So the total energy looks something like

$$E(\{\mathbf{r}_i\}) = \frac{\kappa}{2} \sum_i (|\mathbf{r}_{i+1} - \mathbf{r}_i| - \ell)^2 + \frac{1}{2} \sum_{ij} V(S_i, S_j) u(\mathbf{r}_{i+1} - \mathbf{r}_i) + \frac{1}{2} \sum_{ij} \Delta(\mathbf{r}_{i+1} - \mathbf{r}_i),$$

where the stiffness κ should be large, the function $u(\mathbf{r})$ needs a shape to express the fact that amino acids have their optimal interaction at finite separation of their centers, and $\Delta(\mathbf{r})$ should be relatively short ranged to express the excluded volume effect. We could try to be a

little more realistic and have an extra variable for each amino acid, to keep track of the configuration of the side chain which project from the position \mathbf{r}_i .




FIG. 80 The twenty different amino acids, arranged from most hydrophobic (top left) to most hydrophilic (bottom right). [perhaps should redraw for better consistency with Fig 79; show only R groups?]

Problem 90: Screening. We are assuming that all interactions extend only over short distances, but we also know that there are charged groups. In this problem you'll show that the long ranged Coulomb interaction is screened. For simplicity, let's imagine that everything is happening in an aqueous solution with only two types of ions, one positive and one negative (e.g., a simple salt solution, where the ions are Na^+ and Cl^-). Let the density of the two ions be $\rho_+(\mathbf{x})$ and $\rho_-(\mathbf{x})$, respectively. If the local electrical potential is $\phi(\mathbf{x})$, then in equilibrium the charge densities must obey

$$\rho_{\pm}(\mathbf{x}) = \rho_0 \exp\left[\pm \frac{q_e \phi(\mathbf{x})}{k_B T}\right], \quad (475)$$

where q_e is the charge on the electron and ρ_0 is the density or concentration of ions in the absence of fields. Suppose that we introduce an extra charge Z at the origin. Convince yourself that




FIG. 81 A model for proteins, after Eq (475). Bonds with stiff springs connect neighboring amino acids, which interact through a potential $u(\mathbf{r})$ when they get close. The strength of the interaction is modulated by the identity of the amino acids through the term $V(S_i, S_j)$ in Eq (475).

the potential then obeys

$$\nabla^2 \phi(\mathbf{x}) = \frac{1}{\epsilon} [Zq_e \delta(\mathbf{x}) + q_e [\rho_+(\mathbf{x}) - \rho_-(\mathbf{x})]], \quad (476)$$

where ϵ is the dielectric constant. The combination of these two equations is often called the “Poisson–Boltzmann” model, since Eq (475) is the Boltzmann distribution and Eq (476) is the Poisson equation of electrostatics. [I have avoided issues of units in electrostatics until now .. get this straight, because we need numbers at the end!]

(a.) Show that, if the spatial variations in potential are small, Eq's (475) and (476) can be combined to give

$$\nabla^2 \phi(\mathbf{x}) + \frac{1}{\lambda^2} \phi(\mathbf{x}) = Zq_e \delta(\mathbf{x}). \quad (477)$$

What is the length λ in terms of the other parameters in the problem?

(b.) You may remember that Eq (477) has solutions that decay exponentially far from the origin; this is the same as for a force mediated by the exchange of a massive particle as opposed to the electromagnetic force, mediated by the massless photon.⁶² In this context, Eq (477) is called the Debye–Hückel equation. Solve Eq (477) to give this result explicitly. If the typical concentration of ions in solution is $\rho_0 \sim 100$ mM, what is the value of λ ?

(c.) With only two univalent ion species, their relative concentrations are fixed by neutrality, and thus there is only one parameter ρ_0 that enters the discussion. Generalize the derivation of the linearized Eq (477) to the case where there are many species of ions.

(d.) Going back to the two–species case in Eq (476), can you solve the problem without making the linearizing approximation that leads to Eq (477)? With spherical symmetry it's a one dimensional problem, so at worst you should be able to do this numerically. With ρ_0 in the range of 100 mM as above, how good is the linearized theory?

At the end of all this, does it seem reasonable that even electrostatic interactions are effectively local?

If we set the interaction $V = 0$, Eq (475) describes a polymer that takes a self-avoiding random walk. If $V = -V_0$, then there is a net attraction that causes collapse of the polymer into a more compact phase at low temperature, but this state is still disordered, since there is nothing to prefer one compact configuration over another. If V depends on the amino acid identities, then if we choose the sequence at random the effective interaction between monomers i and j will also be random. Although this sounds like a complicated problem, we know a great deal about the behavior of systems where the Hamiltonian contains terms chosen at random.

⁶² Historically, this idea goes back to Yukawa, who imagined the strong force between protons and neutrons mediated by the exchange of a heavy particle. We now know that this was on the right track, but there were more layers of the strong interaction to be uncovered; solutions to Eq (477) are still called Yukawa potentials. A more direct connection to the standard model of particle physics is in the case of the weak interaction, where the large mass of the W^\pm and Z bosons are directly related to the short range over which the weak interaction is effective.

The prototype of a system with random interactions is the spin glass. Imagine a solid in which, at every site, there is a magnetic dipole which can point up or down, and hence can be described by an Ising spin $\sigma_\mu = \pm 1$ at site μ . If neighboring spins tend to be parallel, then we can write the Hamiltonian as

$$H = -J \sum_{\langle i,j \rangle} \sigma_i \sigma_j, \quad (478)$$

where $\langle i,j \rangle$ denotes neighboring sites. In the classic spin glass materials, magnetic impurities are dissolved in a metal, so the distances between neighbors are random. Further, when the conduction electrons in the metal respond to the magnetic impurity, they polarize, but in a metal all the electronic states involved in responses to small perturbations are near the Fermi surface, and hence have a very limited range of momenta or wavevectors in their wavefunctions. This limitation in momentum space corresponds to an oscillation in real space, so the polarization surrounding a single magnetic impurity oscillates with distance; a neighboring impurity will ‘feel’ this polarization, and so the effective interaction between the two impurities can be positive or negative, at random, depending on the distance between them. This suggests a Hamiltonian of the form

$$H = - \sum_{ij} J_{ij} \sigma_i \sigma_j, \quad (479)$$

where J_{ij} is a random number. In a real system these interactions would be nonzero only for nearby spins, but there is a natural ‘‘mean field’’ limit in which we allow all the spins to interact; this is the Sherrington–Kirkpatrick model.




FIG. 82 Three frustrated spins. Signs on the bonds indicate the signs of J_{ij} in Eq (479). No matter what configuration of spins we choose, one of the bonds is always unsatisfied.

The key qualitative idea in spin glass theory is *frustration*, schematized in Fig 82. In the case of the ‘‘ferromagnetic’’ Ising model in Eq (478), each term in the

Hamiltonian can be made as negative as possible by having all the spins point in the same direction, either up or down. But, in the spin glass case, we may find (for example) that spin 1 is coupled to spins 2 and 3 with ferromagnetic interactions $J_{12} > 0$ and $J_{13} > 0$, but spins 2 and 3 are coupled to each other with an anti-ferromagnetic interaction, $J_{23} < 0$. In such a triangle, there is no configuration of the spins which can optimize all the terms in the energy function simultaneously—the interactions compete. As one can see in this simple problem with three spins, a consequence of this competition is that there are many states of the system with low energy that are nearly degenerate. Importantly, in systems with many spins these low lying states correspond to very different spin configurations.

Problem 91: Simulating (small) spin glasses. Consider a mean field spin glass, as in Eq (479), in which the couplings $J_{\mu\nu}$ are drawn at random from a Gaussian distribution; for simplicity start with the assumption that the mean of this distribution is zero and the variance is one. Notice that with N spins there are exactly 2^N states of the system as a whole, so that up to $N = 20$ (or even a bit more) you can easily enumerate all of these states without taxing the memory of your laptop.

(a.) Write a simple program (e.g., in MATLAB) which, starting from a particular random matrix $J_{\mu\nu}$, gives the energies of all the states in an N spin system.

(b.) Find the ground state energy of an N spin system, and do this many times for independent choices of the random interactions $J_{\mu\nu}$. Show that, if the distribution out of which the $J_{\mu\nu}$ are drawn is held fixed, then the ground state energy does not seem to be extensive (i.e., proportional to N) as N varies. In contrast, if the variance of J scales $\propto 1/N$, show that the average ground state energy does seem to be proportional to the number of spins. Can you give an analytic argument for why this scaling should work?

(c.) The exact ground state energy depends on the particular choice of the interactions $J_{\mu\nu}$. One might hope that, as the system becomes large, there is a ‘‘self-averaging,’’ so that the energy per spin becomes independent of these details in the limit $N \rightarrow \infty$. Do you see any signs of this?

(d.) Having normalized the variance of the couplings $\langle J^2 \rangle = 1/N$, so that the ground state energy is on the order of -1 per spin, compute the gap Δ between the ground state and the first excited state of the system, again for many realizations of the matrix $J_{\mu\nu}$. How does the probability distribution of this gap behave at small values of the gap? In particular, is there a finite probability density as $\Delta \rightarrow 0$? How does this behavior of the gap compare with what you expect in a ferromagnet?

(e.) Show that at least some of the low lying states have spin configurations that are very different from the ground state. Again, contrast this with the case of a ferromagnet.

The statistical mechanics of spin glasses is a very beautiful subject, and we could spend a whole semester on this. What we need for the moment, however, is an intuition, something of the sort one can get from the numerical simulation above. In systems with substantial frustration, we expect that there will be many locally stable

low energy states, and these will be very far apart in the relevant state space. Thus, rather than having a well defined ground state, with small fluctuations around this state, there are many inequivalent near-ground states, often with large barriers between them. If we think of the dynamics of the system as motion on an energy surface, then this surface will be rough, with many valleys separated by high passes; indeed, in the Sherrington–Kirkpatrick model there are valleys within valleys, hierarchically. **This needs a figure. It's a bit conventional, but maybe there is a reason for the convention?**

What does all of this teach us about the protein folding problem? To the extent that we can make analogies between spin glasses and heteropolymers with random sequences, we expect that these randomly chosen proteins will not, in general, have unique ground state structures. Instead, there will be many inequivalent structures with nearly the same low energy, separated by large barriers. Several groups have used modern tools from the statistical mechanics of disordered systems to make this intuition precise **Should I say something about the heftier calculations? An Appendix about replicas? Where else do we really need those ideas?**, and indeed the random heteropolymer is a kind of glass—the polymer has compact, locally stable structures, but there are many of these, and the system tends to get ‘stuck’ in one or another such local minimum at random. This contrasts sharply the ability of real proteins to fold into particular, compact conformations that are (at some level of coarse graining) unique, determined by the sequence. The real problem is even worse, because we have only considered the statistical mechanics of one polymer in solution; in practice the folded state of proteins competes not only with the higher entropy unfolded state, but with states in which multiple protein molecules aggregate and precipitate out of solution.

The conclusion is that the proteins which occur in Nature cannot be typical of sequences chosen at random. At the same time, not every detail of the amino acid sequence can be important. This is perhaps the most fundamental example of the general question we are exploring in this Chapter—our description of life cannot depend on fine tuning, but neither are the phenomena of life generic. Concretely, we can ask how to describe the ensemble of sequences that we see in real proteins. One possibility is that this ensemble is profoundly shaped by history, and surely at some level this is true—we can trace evolutionary relationships through sequence data. Another possibility is that the ensemble of *possible* sequences is enormously constrained by physical principles—ensuring that a protein will fold into some compact, reproducible structure is very difficult, and perhaps even enough to explain the dramatically restricted range of sequences and even structures that we observe in real proteins.

At this point we should pause to note that the prob-




FIG. 83 A schematic energy landscape for protein folding, from Onuchic et al (1995). **[Maybe redraw this? Would be good to have equations in the text to point at for features of the funnel.]**

lem we are formulating is related to, but different from, a much more widely discussed problem. The general question of how protein structure emerges from the underlying amino acid sequence is referred to as the “the protein folding problem.” As a practical matter, one might like to predict the three dimensional structure of the folded state, starting only with the sequence. Many approaches to this problem are based not on a physical model for the interactions, but on attempts to generalize from many known examples of sequence/structure pairs. Faced with a particular sequence from Nature, this is can be an extraordinarily effective approach. But it doesn’t tell us why some heteropolymers fold into compact, reproducible states, while others do not, and why (presumably) some sequences will never be seen in real organisms. It is this more general version of the question that concerns us here.

One approach emphasizes that in a typical sequence chosen at random, interactions among the different amino acids will be frustrated, blocking the system from finding a single well isolated folded structure of minimum energy. A candidate principle for selecting functional sequences is thus the minimization of this frustration. If frustration is absent, there may be few if any major energetic barriers on the path from an unfolded state to the compact, native conformation, although the need for local structural rearrangements along the path may mean that there is an irreducible ‘roughness’ to the energy surface that, in a coarse grained picture, will limit the mobility of the system along its path. This scenario has come to be called a folding ‘funnel,’ emphasizing that there is a single dominant valley in the energy landscape, into which all initial configurations of the system will be drawn, as shown schematically in Fig 83.

At a technical level, if frustration is absent, then we can look at the ground state or native structure and “read off” an approximation to the interactions. Thus, in a ferromagnet, all the spins are parallel in the ground state, and if simply look at each neighboring pair, we would guess that there is a ferromagnetic interaction between them; absent any other data, we should assume that all these interactions have the same strength. Although this might not be exactly right, the Hamiltonian we get in this way will have the correct ground state. In contrast, this doesn’t work with spin glasses, because the (near-)ground states necessarily leave some fraction of the interactions unsatisfied, due to frustration. In this

spirit, if we look at a small protein, we might try to generate a potential energy function which ties neighboring amino acids together along the chain and, in addition, has “bonds” between amino acids which are in contact in the folded state. We should choose the scale of the potential to have more or less the correct distance between amino acids, and the right order of magnitude for the free energy difference between folded and unfolded states.

Models which bond together amino acids that should form contacts, and neglect all other interactions, actually have a long history, and referred to as Gō models. Concretely, this approach involves an energy function of the form

$$E = \frac{1}{2} \sum_{\text{bonds}} \kappa_r (r - r_0)^2 + \frac{1}{2} \sum_{\text{angles}} \kappa_\theta (\theta - \theta_0)^2 + \frac{1}{2} \sum_{\text{dihedrals}} \sum_n \kappa_\phi^{(n)} [1 + \cos(n(\phi - \phi_0))] + \epsilon \sum_{i < j - 3} \left[5 \left(\frac{\sigma_{ij}}{r_{ij}} \right)^1 2 - C_{ij}^{\text{native}} 6 \left(\frac{\sigma_{ij}}{r_{ij}} \right)^{10} \right], \quad (480)$$

where the various κ s are stiffnesses which hold bond lengths r and angles θ, ϕ along the chain to their native values. The crucial terms are those in the second line, which serve to bond together pairs of residues ij which form a contact in the native, folded state ($C_{ij}^{\text{native}} = 1$) while pushing apart those which do not ($C_{ij}^{\text{native}} = 0$). In principle the different bonds can have specific lengths σ_{ij} , but this is not so important qualitatively.

More recently it has been possible to test these ideas in more detail, by complete simulations of the folding process (cf Fig 84). To summarize the results of the simulation, we can measure the fraction Q of the contacts which should form in the folded state that have actually been made; by construction, as this order parameter increases, the energy of the system decreases. But making contacts lowers the entropy of the polymer, and exactly how much the entropy is lowered depends on which contacts are made. When the dust settles, we can see that the free energy as a function Q has roughly a double well structure. Importantly, one can also sample the configurations in the transition state between the wells, and ask which contacts have been made by the time the molecules finds its way to the top of the barrier. Because there are no competing interactions, the prediction is that the ensemble of transition state configurations must reflect only the geometry of the target, folded state.

Can we test the predictions of such simulations? We expect, from the general arguments in Section II.A, that the rate of folding will have an approximately Arrhenius temperature dependence, $k \propto \exp(-\Delta F/k_B T)$, where ΔF is the free energy difference between the unfolded state and the “transition state” at the top of the barrier.




FIG. 84 Gō models for two particular proteins, dihydrofolate reductase (DHFR at left) and interleukin 1β (IL- 1β at right), from Clementi et al (2000). Along the x -axis in all figures is a parameter Q measuring the fraction of native contacts that have formed. The top panels show the root-mean-square difference between the structures and the ground state, with colors denoting the energy. Note that, because there are no competing interactions, the energy decreases linearly as more of the native contacts are formed. But different values of Q can be achieved by different numbers of configurations, until at $Q = 1$ there is only one possible structure. Thus the entropy generally declines with Q , although there is also some structure along the way determined by the geometry of the native fold. The result, shown in the bottom panels, is that the free energy has two distinct minima, corresponding to folded ($Q \approx 1$) and unfolded ($Q \approx 0$) states. Different curves correspond to different temperatures, as indicated.




FIG. 85 Simulations of folding for two proteins, using $G\ddot{o}$ models, from Clementi et al (2000). At each instant of time in the simulation we can count the fraction Q of native contacts, as in Fig 84; sampling the probability distribution of Q we infer the free energy $F(Q)$. At left, simulations of an SH3 domain, which is known to fold rapidly with no obvious intermediate states between folded and unfolded. At right, simulations of the enzyme RNase, which folds more slowly and occupies a well defined intermediate state. These differences are captured by the $G\ddot{o}$ models, suggesting that frustration does not play a role in slowing the folding of the larger molecules.

Imagine that we mutate the protein to change amino acid i . This has some effect on the free energy of every contact between i and j , and we can measure at least the sum of these effects by measuring the change in the free energy difference between the folded and unfolded states. But if along the “reaction coordinate” Q in Fig 84 these contacts are made (on average) only once $Q > Q_c$, where the Q_c is the position of the transition state, then changing their energy doesn’t change the activation free energy for the folding reaction. On the other hand if these contacts are made at $Q < Q_c$, they contribute to the free energy of the transition state and should change the rate of folding. Roughly speaking, the ratio between changes in the (kinetic) free energy of activation and the (thermodynamic) free energy of folding tells us the fraction of contacts involving residue i which are formed in the transition state, and this is something we can get directly from the computations summarized in Fig 84; it is also something one can measure experimentally. Theory and experiment are in surprisingly good agreement [show a figure with the comparison!], which strongly suggests that, at least for small proteins, frustration really has been minimized.

Problem 92: The location of transition states. Suppose that the dynamics of a chemical reaction are described, as in [pointer], by motion of a coordinate x in a potential $V(x)$ that

has two minima separated by a barrier. Let the locations of the two minima be at x_1 and x_2 , while the peak of the barrier is at a position x_t . Assume that rate constants from transitions between the two wells are governed by the Arrhenius law. Now imagine that we apply a small force f directly to the coordinate x . How does this change the equilibrium between the two states? How does it change the rate of transition, say from the states near x_1 to the states near x_2 ? Notice that these are measurable quantities. Can you combine them to infer the location of x_t along the line from x_1 to x_2 ? In particular, can you say something without knowing any additional parameters?

Some proteins are known to fold slowly, moving through a well defined intermediate state. Does this represent a failure to relieve all of the frustration, or is it somehow intrinsic to the size and structure of these molecules? One can make $G\ddot{o}$ models of these slower proteins, and compare them with the smaller “two state folders.” Results of such a comparison are shown in Fig 85. Perhaps surprisingly, intermediates emerge in the folding of the larger protein even in a model where there is no intrinsic frustration from the interactions among different kinds of amino acids. [I’d like to understand if one be more quantitative here ... can we really conclude that frustration is approximately minimized?]

A second approach to our problem looks more explicitly at the mapping between sequences and structures. The observation that changes in amino acid sequence (mutations) don’t necessarily change protein structure tells us that many sequences map into the same structure. But what about the other direction of the mapping? If we imagine some compact structure of a hypothetical protein, can we find a sequence that will fold into this structure? This is the inverse folding problem, or the problem of protein design.




FIG. 86 Compact “folded” structure of an $N = 30$ polymer on a square lattice.




FIG. 87 Exhaustive simulations of compact structures on a lattice, from Li et al (1996). At left, the number of structures which are the ground state for exactly N_s distinct HP sequences, plotted vs N_s for $3 \times 3 \times 3$ (top) and 6×6 (bottom) lattices. Note the small number of structures which are the ground states for huge numbers of sequences. At right, the energy gap between the ground state and the first “excited” state, showing that stability correlates with N_s ; the most highly designable structure has a distinctive pattern of hydrophobic and polar residues alternating with residues that are free to be either H or P with nearly equal probability.

To address the inverse folding problem it is helpful to step back and work on a simpler version of the problem. Imagine that there are just two kinds of amino acids, hydrophobic (H) and polar (P). Polar residues are happy to be next to one another, but they are equally happy to be on the outside surface of the protein, interacting with water. Hydrophobic residues are much happier to be next to one another, and this includes the effect of not being near water. Finally, for hydrophobic residues, it is likely that having a polar neighbor is marginally better than having water as a neighbor. Thus there are three interaction energies, $E_{PP} > E_{HP} > E_{HH}$, where lower energy is (as usual) more favorable. To simplify yet further, let us assume that the structure of the protein lives on a lattice, as in Fig 86. Now it’s clear what we mean by ‘compact’ structures—if the protein is $N = 27$ amino acids long, for example, a compact structure is one which fills a 3×3 cube—and similarly the definition of ‘neighbor’ is unambiguous.

Once we have simplified the problem, it is possible to attack it by exhaustive enumeration. On the $3 \times 3 \times 3$ cube, for example, there are only $\sim 50,000$ inequivalent compact structures, and there are only $2^{27} \sim 10^8$ sequences of this length in the HP model. These numbers are large, but hardly astronomical, so one can explore these sequences and structures completely, also for two dimensional models with $N = 30$ and 36 . To begin, out of 2^{27} sequences, less than 5% have a unique

compact structure with minimum energy; the majority of sequences have multiple degenerate ground states with inequivalent structures. Conversely, there are nearly 10% of compact structures for which no sequence finds that structure as its ground state; the vast majority of structures are connected to just a handful of sequences. But if we ask how many sequences map into a given structure (N_s), there is a long tail to the distribution of this number (Fig 87, at left), and some structures have thousands of sequences that all reach that structure as their ground state. We can say that these structures are easy to design, or ‘highly designable.’ Structures with large N_s also have a large energy gap between the compact ground state and the next highest energy conformation, so that highly designable structures are also thermodynamically stable.

What are these highly designable structures? It is hard to extrapolate from such small systems, but certainly the structures with largest N_s have more symmetry and show hints of extended elements such as helices and sheets, as seen in the insets to Fig 87). Can we understand why designability is so variable, and why these particular structures are highly designable?

Before proceeding it is worth noting that finding sequences that stabilize certain structures can be done in two ways. What we really want are sequences with the property that the desired structure is actually the ground state, which means we have to check all other possible competing structures. A weaker notion is to ask for a sequence that assigns a low energy to the desired structure, perhaps even the lowest possible energy across all sequences. If we are just trying the lower the energy, then the problem of choosing sequences is relatively simple—we should try to put the polar residues on the outside, and the hydrophobic residues on the inside. This version of the inverse problem seems at most weakly frustrated, so there are “downhill” paths to find good sequences. [Is there more to say here?]

Analytic approaches to designability describe protein structure not in terms of the positions of all the amino acids, but in terms of a matrix C_{ij} that specifies whether monomers i and j are in contact ($C_{ij} = 1$) or not ($C_{ij} = 0$); by convention $C_{ii} = 0$. Assuming that all long ranged interactions are screened we can approximate the energy of the molecule as having contributions only from amino acids that are in contact,

$$E = \sum_{ij} C_{ij} \sum_{\mu\nu} s_i^\mu V_{\mu\nu} s_j^\nu, \quad (481)$$

where $s_i^\mu = 1$ if the amino acid at site i is of type μ , and $s_i^\mu = 0$ otherwise. The matrix $V_{\mu\nu}$ summarizes the interactions among the different types of amino acids. To approach the weaker notion of designability, we need to ask how many sequences give a particular structure a low energy. But asking about the numbers of sequences with a particular energy is just like doing statistical mechanics

where we keep the structure fixed and instead allow the sequence $\{s_i^\mu\}$ to be the dynamical variable. This suggests that we compute the partition function in sequence space,

$$Z_{\text{seq}}(C) = \sum_{\{s_i^\mu\}} \exp \left[-\beta \sum_{ij} C_{ij} \sum_{\mu\nu} s_i^\mu V_{\mu\nu} s_j^\nu \right]. \quad (482)$$

Again, this is hard in general, but we can get some intuition by doing a high temperature (small β) expansion.

Summing over all sequences is equivalent to averaging over a distribution in which all sequences are equally likely. Recall that computing the average value of an exponential generates a series of cumulants, or connected correlations:

$$\langle e^{-x} \rangle = \exp \left[-\langle x \rangle + \frac{1}{2} \langle x^2 \rangle_c - \frac{1}{3!} \langle x^3 \rangle_c + \dots \right] \quad (483)$$

$$\langle x^2 \rangle_c = \langle x^2 \rangle - \langle x \rangle^2 = \langle (x - \langle x \rangle)^2 \rangle, \quad (484)$$

$$\langle x^3 \rangle_c = \langle (x - \langle x \rangle)^3 \rangle, \quad (485)$$

and so on. To use this in evaluating $Z_{\text{seq}}(C)$, we need to compute quantities of the form

$$\left\langle \sum_{\mu\nu} s_i^\mu V_{\mu\nu} s_j^\nu \right\rangle,$$

or

$$\left\langle \left(\sum_{\mu\nu} s_i^\mu V_{\mu\nu} s_j^\nu \right)^2 \right\rangle.$$

Since we are averaging over a distribution in which all sequences are equally likely, the vector \vec{s}_i that specifies the choice of amino acid at site i is independent of the vectors \vec{s}_j for any $j \neq i$. Pushing through the details, this allows us to show that the free energy

$$F_{\text{seq}}(C) \equiv -\frac{1}{\beta} \ln Z_{\text{seq}}(C) = A \text{Tr}(C^2) + B \text{Tr}(C^3) + \dots, \quad (486)$$

where the coefficients depend on the details of the potential $V_{\mu\nu}$, and the term $\sim \text{Tr}(C)$ is absent because $\text{Tr}(C) = 0$.

Problem 93: Details of $F_{\text{seq}}(C)$. Derive Eq (486), carrying the expansion out to at least one more order. Relate the coefficients in the expansion explicitly to the properties of the potential $V_{\mu\nu}$.

Because the elements of the matrix C are either 1 or 0, $\text{Tr}(C^2)$ just counts the number of contacts, while $\text{Tr}(C^3)$




FIG. 88 The connection between designability and the eigenvalues of the contact matrix. [explain]. From England & Shakhnovich (2003).

counts the number of connected paths that lead from site i to site j to site k and back to site i . Similarly, the trace of higher powers counts the number of longer paths. But we can also take a less local view and note that $\text{Tr}(C^n) = \sum_i \lambda_i^n$, where λ_i are the eigenvalues of the matrix C . As we consider higher powers in the expansion, the result is dominated more and more by the largest of these eigenvalues. Experimenting with small structures as in the discussion above, one can show that the designability of a structure really does correlate strongly with the largest eigenvalue of the contact matrix, and the most designable structures have the largest eigenvalues, as in Fig 88. This is especially interesting since the calculation we have outlined here does not depend on details of the assumptions about the interactions between amino acids—all that matters is locality.

As noted above, computing $F_{\text{seq}}(C)$ gives us a “weak” notion of designability, counting the number of sequences for which a particular structure will have low energy. If we are willing to simplify our model of the interactions, then we can make progress on the stronger notion of designability, that many sequences have the same minimum energy structure. Suppose we return to the model in which there are just two kinds of amino acids, hydrophobic and polar. Further, let’s describe the structure in a similar binary fashion, labeling each amino acid by whether it is on the surface of the molecule or in the interior.⁶³ Now there is a plausible energy function—hydrophobic residues prefer interior sites, polar residues prefer the surface. Thus the energy will be minimized

⁶³ On a lattice, with the protein folded into a compact structure, this categorization of sites is unambiguous, although one might worry a bit about the more general case.

when the binary description of the sequences ($s_i = +1$ for hydrophobic, $s_i = -1$ for polar) matches the binary description of the structure ($\sigma_i = +1$ for interior, $\sigma_i = -1$ for the surface). Although we might not be able to calculate the exact energy function, ground state structures should correspond to the minimum of a very simple energy that just counts the violations of the hydrophobic/interior, polar/surface rule,

$$E \propto \sum_i (s_i - \sigma_i)^2. \quad (487)$$

An important point about this binary description of structures and sequences is that while all binary strings $\{s_i\}$ represent possible amino acid sequences, not all binary strings $\{\sigma_i\}$ are possible compact structures of a polymer [maybe it would be useful to have a figure illustrating this point?]. Thus in the space of binary strings, and hence H/P sequences, there are special points that correspond to realizable protein structures. The energy function in Eq (487) tells us that the ground state structure for any sequence is the nearest such point, where “near” is measured by a natural metric, the “Hamming distance,” counting the number of bits that disagree in the binary string. The set of sequences that will fold into one particular structure are those which fall within the Voronoi polygon surrounding the binary description of that structure, as shown in Fig 89. In this picture, the sequence literally encodes the structure, and the folding process provides a kind of error correction in this code, mapping arbitrary binary strings back to the sparse set of realizable structures. By choosing structures which are far from other structures in this binary representation, one guarantees that many sequences will map to that one structure. Again this picture can be tested against simulations of the lattice models (as in the discussion above), and the results are consistent.

The lesson from all this is that not all structures are created equal, and that selection of structures for their designability induces a nontrivial distribution on the space of sequences. This constraint of course restricts the set of allowed sequences, but at the same time focuses precisely on those sequences for which not all details of the sequence have functional relevance. [check if there is more worth saying here]

There is yet another approach which tries to address the ensemble of allowed sequences, leaning on theory but also using a more direct experimental exploration. In order to appreciate this approach, you need to know that proteins form families. We have already met a simple example of this, with rhodopsin. In your retina, there are four kinds of photoreceptor cells—rods for night vision, and three kinds of cones that provide color vision at higher light intensities—and each one expresses a different pigment molecule, with a different absorption spectrum. Rhodopsin consists of a medium sized organic

molecule, retinal embedded in the protein; all the pigments use retinal, so the differences in absorption spectrum reflect differences in the protein. All of these proteins are doing the same job, and have recognizably related structures and amino acid sequences. Nonetheless, they are not identical. In fact, they share sequence and structural similarities with many more proteins, all of which function as receptors (usually for the binding of small molecules rather than the absorption of light), and sit in a membrane rather than being free in solution. Rhodopsin interacts with transducin (Section I.C), which functions as the first stage of an amplification cascade, and other rhodopsin-like molecules interact with similar amplifier molecule. The family to which transducin belongs is called the “G proteins,” because part of their function is driven by the hydrolysis of GTP to GDP [be sure this was clear in Chapter 1!], while the rhodopsins and relatives are referred to as G protein coupled receptors (GPCRs). There are GPCRs that respond to hormones, to neurotransmitters in the brain, and, notably, to odorants in the receptor cells of the nose.

Important examples of protein families are provided by enzymes. For example, there are many enzymes which attach phosphate groups to other proteins, for example, and there is variety even within an organism because these protein kinases have different targets; there is even more diversity across organisms. In order to digest our food, we need to cut up the proteins that we ingest, and all cells also need to cut up old proteins that have been damaged or outlived their usefulness in other ways. Cutting the peptide bond quickly and efficiently requires a carefully engineered catalyst, but cells also need control over which sequences they are cutting. Thus there are several families of protein-cutting proteins, called proteases, and there are remarkable structural similarities among molecules separated by billions of years of evolu-




FIG. 89 Designability as seen in the binary description of sequences and structures. [explain]. From Li et al (1998).




FIG. 90 Comparison of the structure of SGPA (right) and chymotrypsin (left), in the neighborhood of the active site; from Brayer et al (1978). Note in particular the very similar geometrical relations among His57, Asp 102 and Ser 195, the triad of residues involved in the catalytic events.

tionary history. An example is shown in Fig 90, comparing the structure of the bacterial enzyme SGPA and the mammalian enzyme chymotrypsin. These molecules have recognizably similar amino acids along only $\sim 25\%$ of their sequences, yet the structures are very similar, especially in the active site where the crucial chemical events occur—the proteins fold to bring these key elements into a very specific geometrical arrangement, despite the sequence differences. Other interesting examples of protein families include smaller parts of proteins (domains) which can fold on their own and function as the interfaces between different molecules; there are hundreds of examples in some of these families.

If we line up the sequences for all the proteins in a family,⁶⁴ as in Fig 91, we find that, at each site there are some preferences for one amino acid over another. With enough members in the family, we get a decent estimate of the probability that an amino acid will be chosen in each position along the sequence. Perhaps the simplest hypothesis about the ensemble of allowed sequences is that amino acids are chosen independently at every site, with these probabilities. It should be emphasized that such ‘one body’ constraints are strong, reducing the entropy of the allowed sequences from a nominal $\sim \log(20)$ per site down to $\sim \log(3)$ per site [check the exact num-

⁶⁴ We need to explain that sequence alignment is not trivial. One might even note that algorithms for alignment (or for the recognition of family members) already embody hypotheses about the answer to the question we are trying to formulate here. This all needs some discussion, not least because it points to open problems!

bers!]. But, this is not enough: if we synthesize proteins at random out of this distribution, it is almost impossible to find one which folds into something like the functional structure characteristic of the original family.

Given that one body models don’t work, it seems the next logical step is to look at two body effects: looking across the family of proteins, we see that substitutions at one site tend to be correlated with substitutions at other sites. Can we sample an ensemble of sequences that captures these pairwise correlations? Let us imagine, for simplicity, that there are only two kinds of amino acid; the real case of twenty possibilities just needs more notation. Then we can use $\sigma_i = +1$ for one kind of amino acid at position i , and $\sigma_i = -1$ for the other. The relative frequency of the two choices is measured by the




FIG. 91 Alignment of the WW domains, showing (A) the sequences in the family and (B) the correlations between amino acids at pairs of sites, measured by the mutual information. The amino acids are indicated by the one letter codes from Fig 80, with – for gaps. Figure from Mora & Bialek (2011), based on data from [explain source!].

“magnetization” $\langle \sigma_i \rangle_{\text{expt}}$, where the subscript remind us that we measure this from data. Similarly, the correlations between amino acid substitutions at pairs of sites is measured by

$$C_{ij}^{\text{expt}} \equiv \langle \sigma_i \sigma_j \rangle_{\text{expt}} - \langle \sigma_i \rangle_{\text{expt}} \langle \sigma_j \rangle_{\text{expt}}. \quad (488)$$

Imagine creating an artificial family of M sequences $\{\sigma_i^\mu\}$, with $\mu = 1, 2, \dots, M$. From this set of replica sequences we can compute the same expectation values that we computed from the real family of sequences,

$$\langle \sigma_i \rangle_{\text{model}} = \frac{1}{M} \sum_{\mu=1}^M \sigma_i^\mu \quad (489)$$

$$C_{ij}^{\text{model}} = \frac{1}{M} \sum_{\mu=1}^M \sigma_i^\mu \sigma_j^\mu - \langle \sigma_i \rangle_{\text{model}} \langle \sigma_j \rangle_{\text{model}}. \quad (490)$$

We would like to arrange for the model family of sequences to have these quantities match the experimental ones. The first part ($\langle \sigma_i \rangle_{\text{model}} = \langle \sigma_i \rangle_{\text{expt}}$) is easy, since we can do this just by choosing the amino acids at every site independently with the same probabilities as in the experimental family. For the two-point correlations, we can form a measure of error between our model sequence ensemble and the real family,

$$\chi^2 = \sum_{ij} \left| C_{ij}^{\text{model}} - C_{ij}^{\text{expt}} \right|^2, \quad (491)$$

and then we can promote this mean square error to an energy function, and adjust the M sequences according to a Monte Carlo simulation with slowly decreasing (effective) temperature. At low temperatures, this procedure should generate an ensemble of sequences which reproduce the pairwise correlations in the naturally occurring sequences. This procedure has been implemented for a real family of proteins, and novel sequences drawn out of the resulting ensemble have been synthesized. Remarkably, a finite fraction of these sequences fold into something close to the proper native structure, and these folded states are essentially as stable as are the natural proteins. [Reproduce a figure from the Ranganathan work?]

In the limit that we are considering a very large family ($M \rightarrow \infty$) of artificial sequences, and we really take the effective temperature to zero, the Monte Carlo procedure draws samples out of a probability distribution that perfectly matches the measured one-point and two-point correlations, but otherwise is as random or unstructured as possible, and hence has maximum entropy. We will meet the maximum entropy idea again in Section III.D, with more details in Appendix A.8. For now, we note that the maximum entropy distribution of sequences takes the form

$$P(\{s_i\}) = \frac{1}{Z} \exp \left[\sum_{i=1}^N u_i(s_i) + \frac{1}{2} \sum_{i,j=1}^N \mathcal{V}_{ij}(s_i, s_j) \right], \quad (492)$$

where the “fields” u_i and the “interactions” \mathcal{V}_{ij} must be chosen to reproduce the one-point and two-point correlations, where now we allow for the amino acid identity at each site to take on all twenty values, $s_i = 1, 2, \dots, 20$. Actually finding these fields and interactions is the inverse of the usual problem in statistical mechanics, and can be challenging. But if we can solve this problem, the maximum entropy method provides a potential answer to the question we posed at the outset—if random sequences don’t fold, and the exact sequence doesn’t matter, how do we describe the ensemble of sequences consistent with a given protein structure or function? Equation (492) gives an explicit answer, a formula for the probability that a particular sequence will occur. Importantly, the form of the distribution is the same as the Boltzmann distribution, with the interactions and fields defining an effective energy surface on the space of sequences. [not sure how to end this .. maybe depends on what Thierry finds in reanalysis of WW domains]

Problem 94: A small maximum entropy model. Give a problem that takes the student through the maxent problem for three spins. Emphasize distinction between interaction and correlation—how much correlation can you get without any direct interactions?

We recall from other problems in statistical mechanics that correlations can extend over much longer distances than the underlying interactions. Thus, although we may detect significant correlations among the amino acid substitutions at many pairs of sites, it is possible that these can be explained by Eq (492) with the interactions \mathcal{V}_{ij} being nonzero only for a very small fraction of pairs ij . Since the physical interactions between amino acids are short ranged, it seems reasonable that if there is a direct connection between the joint choice of residues at sites i and j on the probability that the resulting protein is a member of the family, then sites i and j should be physically close to one another in the protein structure. This idea was worked out in detail for pairs of receptors and associated signaling proteins in bacteria, and it was possible to identify, with high reliability, the amino acids which make up the region of contact between these molecules, as shown in Fig 92. This success raises the tantalizing possibility that we could read off the physical contacts between amino acids—and hence infer the three-dimensional structure of proteins—from analysis of the covariations in amino acid substitutions across a large family.

Should end with some review of what we have learned about the interplay of tuning and robustness; at least some of these questions have become more quantitative.




FIG. 92 Interactions between residues in the ensemble of sequences predict spatial proximity, from Weigt et al (2009). [Fill in caption! Do we need more discussion in the text to define “direct information” as generalization of J_{ij} ?

There is also a question about history vs. physics: is the ensemble of sequences *just* a record of evolutionary history, or more like an equilibrium distribution subject to some sensible physical constraints? Do we want to say something explicit about the antibodies? Emphasize that the challenge of building the maximum entropy distributions for larger proteins is really still open?

The amino acid sequences of proteins are translations of the DNA sequences. But there are large parts of DNA which are not coding for proteins. Important parts of this “non-coding” DNA are involved in transcriptional regulation, as discussed in Section II.B. The key steps of this regulatory process involve the binding of transcription factor proteins to DNA, and the architecture of the regulatory network depends on the specificity of these protein–DNA interactions. When we draw an arrow from one transcription factor (TF) to its target gene, then as schematized in Fig [** we had a schematic in a previous chapter, but maybe need another one here?] there must be a short sequence of DNA in or around the target gene to which the transcription factor can bind. The fact that a given TF activates or represses one gene, but not another, then is controlled by the presence or absence of the relevant sequences. But some transcription factors are quite promiscuous, and in higher organisms the relevant sequences often are quite short, so this specificity is not all-or-none. Rather we should think that every short sequence is a possible binding site, and there is a binding energy that depends on the sequence.

Formally, a short piece of DNA sequence can be thought of as a series of bases. Let’s write $s_i^\mu = 1$ if the base at position i is of type μ ; we have $\mu = 1, 2, 3, 4$ and $i = 1, 2, \dots, L$, where L is the length of the possible binding site. We can abbreviate $\mathbf{s} \equiv \{s_i^\mu\}$. Then if we

look at one transcription factor, there is some binding energy of that factor to the DNA, $E(\mathbf{s})$, for every possible sequence. What does the function $E(\mathbf{s})$ look like? Obviously, if it’s a constant then there is no specificity at all—a given transcription factor will influence every gene in the genome—and this can’t be right. On the other hand, if the binding is strong only for one specific sequence \mathbf{s}_0 (that is, $E(\mathbf{s}) = -E_0$ with large $E_0 > 0$), while $E(\mathbf{s} \neq \mathbf{s}_0) \sim 0$, then the transcription factor can successfully target a small subset of genes, but the landscape for evolutionary change becomes very rugged—changing a single base can completely eliminate one of the regulatory “arrows” in the network, or create a new one of equal strength to all previous arrows—and this doesn’t seem right either.

We can turn our question about the form of $E(\mathbf{s})$ around and ask about the set of sequences that will act as functional binding sites, presumably those sequences that have $E(\mathbf{s})$ in some range. In one limit, this ensemble would include all sequences; in the other limit, there would be just one sequence. Thus the issue of specificity in protein–DNA interaction is rather like the problem of amino acid sequence ensembles with which we started this Chapter: where do real biological systems sit along the continuum between completely random sequences at one extreme and unique sequences at the other?

Many of the ideas for analyzing the nature of the sequence ensemble for binding sites involve the starting assumption that each base contributes linearly to the total binding energy, so that

$$E(\mathbf{s}) = \sum_{i=1}^L \sum_{\mu=1}^4 W_{i\mu} s_i^\mu, \quad (493)$$

where $W_{i\mu}$ are the weights given to each position i . One of the first ideas was, in the language we have already used, a maximum entropy argument. If all we know is that functional binding sites must have some average binding energy $\langle E \rangle$, then the maximum entropy distribution consistent with this knowledge is

$$P(\mathbf{s}) = \frac{1}{Z} \exp[-\lambda E(\mathbf{s})], \quad (494)$$

which of course is the Boltzmann distribution at some effective temperature $\propto 1/\lambda$. Importantly, if the energy is additive as in Eq (493), then the probability of the entire sequence is a product of probabilities at the different sites,

$$P(\mathbf{s}) = \frac{1}{Z} \prod_{i=1}^L \exp \left[-\lambda \sum_{\mu=1}^4 W_{i\mu} s_i^\mu \right]. \quad (495)$$

This means that the expected frequency of occurrence of the different bases at each site—that is, the probability that $s_i^\mu = 1$ —can be related directly to the weight matrix,

$$f_{i\mu} \propto \exp[-\lambda W_{i\mu}]. \quad (496)$$

Thus, if we could get a fair sampling of the ensemble of sequences we could just read off the matrix elements $W_{i\mu}$. [Should I explain that Berg & von Hippel never said “maximum entropy”? Does it matter?]

Problem 95: Random sequences. Take the students through expectations about the distribution of binding energies for the case where sequences are random.

When these ideas first emerged in the mid to late 1980s, in work by Berg & von Hippel, there were few examples where one could point to multiple known binding sites for a single transcription factor. Two important examples were the *lac* operon and the phage λ switch. These are sufficiently important examples in the history of the subject that it is worth taking some time to explain here how they work. [Do this!]

Problem 96: A little more about λ . Depends on what gets said in the text, but maybe ask the students to reproduce Ptashne’s argument about the importance of cooperativity.

What was available to Berg and von Hippel were ~ 100 examples of the DNA sequences to which RNA polymerase binds when it begins transcribing. This of course is another example of protein–DNA interaction, not a regulatory interaction but an essential part of all gene expression.⁶⁵ Further, there had been in vitro kinetic measurements on transcription, so they knew something about directly about the binding energies. If experiments are done in the regime where the binding sites are usually empty, then the observed transcription rates will be proportional to the concentration of polymerase and the equilibrium constant $K \propto \exp[-\beta E(s)]$. The comparison is shown in Fig 93, including some estimates of errors in the measurements and predictions. The agreement is quite good. Thus, it really does seem that one can, at least roughly, estimate the energetics of binding events

⁶⁵ Even in this case the number of sequences is not very large, and we should remember that we are trying to estimate the frequencies of four different bases at each site. To improve their estimates, Berg & von Hippel (1987) used “psuedo-counts,” a procedure explained in Appendix A.9.

from the statistics of sequences, which is quite surprising.

The sequencing of whole genomes, from many organisms, created the opportunity for much more systematic exploration of sequence ensembles. The fact that the number of transcription factors is very much smaller than the number of genes means that, generally, even in a single organism there must be many examples of binding sites for each transcription factor. It seems likely, then, that similar sequences—sequences with good binding energies—will appear more frequently than would be expected at random, and these sequences should, in the simplest cases, be positioned near the start sites of transcription.

In written language, short sequences of letters that occur more frequently than expected by chance have a name—words. When we read, however, there are spaces and punctuation that mark the limits of the words, so we can recognize them. Interestingly, this is less true for spoken language, where the sounds of words often run together, and pauses or gaps are both less distinguishable and less reliable indicators of word boundaries. In fact, we really don’t need these markers, even in the case of written text, as you can see by reading Fig 94.

In the simplest view, words are independent, and all structure arises from the fact that not all combinations of




FIG. 93 Sequence dependence of RNA polymerase activity compared with predictions from a maximum entropy model, from Berg & von Hippel (1987). On the vertical axis, effective second-order rate constants for the initiation of transcription by combination of RNA polymerase and different promoter sequences. On the horizontal axis, scaled binding energies predicted from a maximum entropy model based on ~ 100 sequences. Points refer to independent biochemical experiments, with lines connecting measurements on the same sequences, giving a sense for the error bars. A solid line with slope -1 is shown to guide the eye, with dashed lines indicating roughly the errors in the model arising from the finite sample size.

theresmanalloverforyou
 blamingonhisbootsthefa
 ultsofhisfeethisisgetting
 alarmingoneofthethieves
 wassaveditsareasonable
 percentagegogo

FIG. 94 A passage from Beckett’s *Waiting for Godot*, spoken by Vladimir. All punctuation and spaces have been removed, but (hopefully) the text can still be understood.

letters form legal words. Then, if we know the boundaries between words, the probability of observing a particular text becomes

$$P = \prod_w [P(w)]^{n_w}, \quad (497)$$

where n_w is the number of occurrences of the word w in the text, and $P(w)$ is the probability of this word. But we don’t really know, a priori, the correct way of segmenting the text into words, and so we need to sum over all possible segmentations. Each segmentation S generates a different combination of words, so the count $n_w(S)$ depends on S . On the other hand, the probability that a word appears is a property of the language, not of our segmentation, and should be constant. Then

$$P = \sum_S \prod_w [P(w)]^{n_w(S)}. \quad (498)$$

If we think of this as a model for a long text, then given the vocabulary defined by the set of possible words $\{w\}$, maximizing the likelihood of the data amounts to setting the predicted probability of each word to the mean number of occurrences of that word when averaged over all segmentations. Because the text is one-dimensional, there are methods to sum over segmentations that are analogous to transfer function methods for one-dimensional models in statistical mechanics. The real challenge in looking at a genome is that we don’t know the vocabulary.

One approach to learning the vocabulary is iterative: start with the assumption that words are single letters, then add two letter words when the frequency of letter pairs is significantly higher than predicted by the model, and so on. To capture the the functional behavior of real biological systems one needs to include words with

gaps, such as TTTCCNNNNNNNGGAAA, in which “N” can be any nucleotide. Indeed, this example is one of the longer words that emerges from an analysis of possible regulatory regions of the yeast genome, and corresponds to the binding site for MCM1, a protein involved in (among other things) control of the cell cycle. Globally, this approach to “building a dictionary” identifies hundreds of words of more than four bases that pass reasonable tests of significance. At the time of the original work, there were ~ 400 known, non-redundant binding sites whose function had been confirmed directly by experiment, and the dictionary reproduced one quarter of these, a success rate 18 standard deviations outside what might have been expected by chance.⁶⁶ One can do even better by repeating the analysis using as input text only the regulatory regions of genes whose expression level is affected during particular processes or by the deletion or over-expression of other genes. More power is added to the analysis by using the genomes of closely related organisms. [What do we want to conclude from all of this? Have we lost the notion of binding energy in this discussion?]

Problem 97: Summing over segmentations. Give a problem to connect summing over segmentations with transfer matrix. See Bussemaker et al (2000b).

A very different approach to our problem involves exploring sequence space more systematically. In a relatively short time, several different technologies have emerged for doing this, each of course with its own strengths and weaknesses. [Explain protein binding microarrays, methods from the Quake lab for similar binding measurements, ChIP methods (but chip and seq). Need one good figure illustrating all of these schematically!! Justin provided some input that I haven’t digested yet here!]

How do we analyze all these data? Certainly we have the impression that this new generation of experiments provides much more systematic, quantitative data, but there are problems. In the protein binding microarray, for example, there seem to be no reliable calibration of the relation between fluorescence levels and binding probability. Certainly if we see a very bright spot, we can be sure that the protein is bound, but the actual distribution of fluorescence intensities has a long tail, as in Fig 95. Where in this tail do we decide that we have a “hit”?

⁶⁶ Say something about what chance means here, and about the general problem of statistical significance in bioinformatics




FIG. 95 Protein binding microarray data on the yeast transcription factor Abf1, from Kinney et al (2007). In blue, a histogram of the fluorescence intensities (relative to background) across all ~ 6000 regulatory regions from the yeast genome (Murkherjee et al 2004). In green, the line drawn in the original experiments to define the threshold for binding. In red, with error bars, estimates of the probability that binding has occurred as a function of the fluorescence level, from the analysis described in the text.

In the experiments of Fig 95, fluorescence is a proxy for protein binding, and if things come to equilibrium then this depends on the DNA sequence through the binding energy $E(\mathbf{s})$. The space of sequences is huge, but the model of Eq (493) says that the binding energy is a linear function of the sequence. Thus, fluorescence should depend on sequence only through a single linear projection. Finding this projection is an example of the dimensionality reduction problem discussed in Appendix A.7. The key idea is that, no matter how complicated or noisy the relationship that connects energy to binding to fluorescence, the sequence can't provide more information about the output of the experiment than it does about the more fundamental quantity $E(\mathbf{s})$. Similarly, if we try to summarize the sequence by any reduced description, we will lose information unless our reduction corresponds to estimating $E(\mathbf{s})$ itself. Thus, if we search for a one dimensional description, corresponding to a single linear projections of the sequence that preserves as much information⁶⁷ as possible about the experimental output, then the projection we find must be our best linear approximation to $E(\mathbf{s})$, up to a scale factor.⁶⁸

Figure 96 show examples of the weight matrices $W_{i\mu}$

⁶⁷ “Information” here is used in the technical sense, in bits. See Section IV.A.

⁶⁸ The actual computation is a bit more involved because the possible regulatory regions are much larger than the binding sites, and so we have to test not all projections, but all possible projections along the relevant ~ 500 base regions. For details see Kinney et al (2007).

obtained from the “maximally informative dimension” analysis of experiments on the yeast transcription factor Abf1, which is assumed to interact with a 20 base long segment of the DNA. Individual matrix elements typically are determined with better than 10% accuracy, and the interaction of the protein with the DNA evidently is dominated by two approximately symmetric regions of five bases, separated by a gap of another five bases. Importantly, using this method it is possible to analyze in vitro (protein binding microarray) and in vivo (ChiP) experiments, and get consistent answers. In contrast, if we just draw a conservative threshold on the signals strengths (e.g., the green line in Fig 95), then these different sorts of experiments typically lead to divergent interpretations. Once we have confidence in the estimates of $E(\mathbf{s})$, we can go back and ask how the probability that the protein is bound is related to the fluorescence intensity, and this is shown in Fig 95. There is nothing about the analysis that forces this relationship to be smooth or monotonic, but it is.

Can we go further, and relate these linear models of binding energy to the control of gene expression itself? Suppose that we put the expression of a fluorescent protein under the control of a known promoter, and then randomly mutate the sequence. We can then generate an ensemble of bacteria with slightly different sequences, each of which will express the fluorescent protein at different levels, presumably because the relevant transcription factor is binding more or less strongly. Experiment-




FIG. 96 Weight matrices $W_{i\mu}$ for Abf1 in yeast, from analysis of ChiP (top) and protein binding microarray (bottom) experiments (Kinney et al 2007). In these analyses the overall scale of $E(\mathbf{s})$ is not determined by the data, and so the two results have been scaled to maximize their similarity. Importantly, the two experiments are done in vivo and in vitro, respectively, but nonetheless generate very similar estimates of the underlying matrix governing protein–DNA interactions. The two matrix elements with the poorest agreement are circled, but even these differences have little effect on the predicted binding energies.




FIG. 97 Analysis of experiments in which the expression of a fluorescent protein is placed under the control of promoter sequences that are randomly mutated versions of the native sequence binding the transcription factor CRP, from Kinney et al (2010). At the top, separate analyses yield the weight matrices $W_{i\mu}$ for the CRP binding site and for the RNA polymerase binding site, up to an arbitrary scale factor. At bottom, a combined analysis places these energies on an absolute scale and determines the interaction energy ϵ_i .

tally, one can sort the cells by their fluorescence, and sequence the promoter regions, and then search once more for a reduction of dimensionality that captures as much information as possible. If the mutations are sprinkled throughout the promoter region, we expect that there are at least two relevant dimensions, corresponding to the binding energy of the transcription factor and the binding energy of the RNA polymerase. The results of such an experiment and analysis are shown in Fig 97.

As before, the search for maximally informative dimensions does not determine the scale of the energies. But if we take seriously that the quantities emerging from the analysis really are energies, then we should be able to compute the probability that the RNA polymerase site is occupied, and it is this occupancy that presumably controls the initiation of transcription. If the energies for binding of the transcription factor (CRP) and RNA polymerase are ϵ_c and ϵ_r , respectively, then the probability of the polymerase site being occupied is

$$\tau = \frac{1}{Z} C_r e^{-\epsilon_r/k_B T} \left(1 + C_c e^{-\epsilon_c/k_B T} e^{-\epsilon_i/k_B T} \right), \quad (499)$$

where the partition function

$$Z = 1 + C_c e^{-\epsilon_c/k_B T} + C_r e^{-\epsilon_r/k_B T} + C_r C_c e^{-\epsilon_r/k_B T} e^{-\epsilon_c/k_B T} e^{-\epsilon_i/k_B T}, \quad (500)$$

where C_c and C_r are the concentrations of the transcription factor and the RNA polymerase, and ϵ_i is the interaction energy between the two proteins when they are both bound to the DNA. Notice that the two binding

energies are quantities whose relation to the sequence should already have been determined by search for maximally informative dimensions, except for the scale and zero of energy. By trying to combine these energies we need to set the scale ($k_B T$) and the zero (equivalently, the concentrations of the proteins), and we have to fit one more parameter, the interaction energy ϵ_i . All of this works, with the results shown at the bottom of Fig 97. For this particular system there are independent measurements of ϵ_i , and there is agreement with $\sim 10\%$ accuracy. Even better, one can show that the single number τ is Eq (499) captures as much information about the sequence dependence of the expression level as do the two numbers ϵ_c and ϵ_r . All of this gives us confidence that the use of statistical mechanics and linear energy models really does make sense here.

Problem 98: RNA polymerase occupancy. Derive Eq (499). Generalize to the case where there are two or more transcription factors, each of which can “touch” the RNA polymerase and contribute an interaction energy. Show that even if the binding of each transcription factor is independent (that is, there are no direct interactions among the TFs), their mutual interactions with the RNA polymerase gives rise to an effective cooperativity in the regulation of transcription. What is the relation of this picture to the MWC models of cooperativity discussed in Appendix A.4?

Now that we have some confidence in our description of the binding energies, we can go back and ask once more about the statistics of sequences, and problem of robustness vs fine tuning. **There are several things to say here. I'd like to cover what happens in Sengupta et al (2002) and Mustonen et al (2008). I think that Justin's observation that you can't find a linear model which points to random collections of genes also is interesting. I'm a bit worried that all of this discussion is in the context of single celled organisms, but there is a lot of stuff to say, e.g., about flies. This needs ALOT of work.**

A good general reference about proteins is Fersht (1998). For a modern introduction to polymer physics, see de Gennes (1979). The small simulation in the problems is not a substitute for exploring the theory of spin glasses; the classic papers are collected, with an introduction, by Mézard et al (1986), and a textbook account is given by De Dominicis & Giardina (2006). Early efforts to apply these methods to the random heteropolymer were made by Shakhnovich & Gutin (1989).

De Dominicis & Giardina 2006: *Random Fields and Spin Glasses* C De Dominicis & I Giardina (Cambridge University Press, Cambridge, 2006).

Fersht 1998: *Structure and Mechanism in Protein Science: A Guide to Enzyme Catalysis and Protein Folding* AR Fersht (WH Freeman, San Francisco, 1998).

- de Gennes 1979:** *Scaling Concepts in Polymer Physics* PG de Gennes (Cornell University Press, Ithaca, 1979).
- Mézard et al 1986:** *Spin Glass Theory and Beyond* M Mézard, G Parisi & MA Virasoro (World Scientific, Singapore, 1986).
- Shakhnovich & Gutin 1989:** Formation of unique structure in polypeptide chains: Theoretical investigation with the aid of a replica approach. EI Shakhnovich & AM Gutin, *Biophys Chem* **34**, 187–199 (1989).
- Models which incorporate only native interactions, with no frustration, have their origin in work by Gō, reviewed in Gō (1983). A more explicit discussion of minimizing frustration as a principle was given by Bryngelson & Wolynes (1987), and the funnel landscape of Fig 83 is from Onuchic et al (1995). Detailed simulations based on the Gō model are described by Clementi et al (2000a,b).
- Bryngelson & Wolynes 1987:** Spin glasses and the statistical mechanics of protein folding. JD Bryngelson & PG Wolynes, *Proc Nat'l Acad Sci (USA)* **84**, 7524–7528 (1987).
- Clementi et al 2000a:** How native-state topology affects the folding of dihydrofolate reductase and interleukin-1 β . C Clementi, PA Jennings & JN Onuchic, *Proc Nat'l Acad Sci (USA)* **97**, 5871–5876 (2000).
- Clementi et al 2000b:** Topological and energetic factors: What determines the structural details of the transition state ensemble and “en-route” intermediates for protein folding? An investigation for small globular proteins. C Clementi, H Nymeyer & JN Onuchic, *J Mol Biol* **298**, 937–953 (2000).
- Gō 1983:** Theoretical studies of protein folding. N Gō, *Ann Rev Biophys Bioeng* **12**, 183–210 (1983).
- Onuchic et al 1995:** Toward an outline of the topography of a realistic protein-folding funnel. JN Onuchic, PG Wolynes, Z Luthey-Schultern & ND Socci, *Proc Nat'l Acad Sci (USA)* **92**, 3626–3630 (1995).
- The lattice simulations which explored protein designability were by Li et al (1996). The analytic argument connecting designability to the eigenvalues of the contact matrix was given by England & Shakhnovich (2003), and Li et al (1998) gave the argument relating folding to error correction in the HP model. [\[Probably there is more to say here!\]](#)
- England & Shakhnovich 2003:** Structural determinant of protein designability. JL England & EI Shakhnovich, *Phys Rev Lett* **90**, 218101 (2003).
- Li et al 1996:** Emergence of preferred structures in a simple model of protein folding. H Li, R Helling, C Tang & N Wingreen, *Science* **273**, 666–669 (1996).
- Li et al 1998:** Are protein folds atypical? H Li, C Tang & NS Wingreen, *Proc Nat'l Acad Sci (USA)* **95**, 4987–4990 (1998).
- [\[Need to start with a general reference about protein families\]](#)
The idea of protein families was essential in the experiments that searched for, and found, the receptors in the olfactory system (Buck & Axel 1991, Axel 2005, Buck 2005). [\[should give general reference for serine proteases\]](#) The structural correspondence between bacterial serine proteases and their mammalian counterparts is from Brayer et al (1978, 1979) and Fujinaga et al (1985). Experiments on the sampling of sequence space while preserving one-point and two-point correlations were done by Socolich et al (2005) and by Russ et al (2005). The equivalence of these ideas to the maximum entropy method was shown in Bialek & Ranganathan (2007). For more on maximum entropy approaches to sequence ensembles, see Weigt et al (2009), Halabi et al (2009), and Mora et al (2010). For a broader view of maximum entropy models applied to biological systems, see Appendix A.8 and Mora & Bialek (2011).
- Axel 2005:** Scents and sensibility: A molecular logic of olfactory perception. R Axel, in *Les Prix Nobel 2004*, T Frängsmyr, ed, pp 234–256 (Nobel Foundation, Stockholm, 2004).
- Bialek & Ranganathan 2007:** Rediscovering the power of pairwise interactions. W Bialek & R Ranganathan, arXiv:0712.4397 [q-bio.QM] (2007).
- Brayer et al 1979:** Molecular structure of crystalline *Streptomyces gresius* protease A at 2.8 Å resolution: II. Molecular conformation, comparison with α -chymotrypsin, and active-site geometry. GD Brayer, LTJ Delbaere & MNG James, *J Mol Biol* **124**, 261–283 (1978).
- Brayer et al 1979:** Molecular structure of the α -lytic protease from *Mycobacter 495* at 2.8 Å resolution. GD Brayer, LTJ Delbaere & MNG James, *J Mol Biol* **131**, 743–775 (1979).
- Buck & Axel 1991:** A novel multigene family may encode odorant receptors: A molecular basis for odor recognition. L Buck & R Axel, *Cell* **65**, 175–187 (1991).
- Buck 2005:** Unraveling the sense of smell. LB Buck, in *Les Prix Nobel 2004*, T Frängsmyr, ed, pp 267–283 (Nobel Foundation, Stockholm, 2004).
- Fujinaga et al 1985:** Refined structure of α -lytic protease at 1.7 Å resolution: Analysis of hydrogen bonding and solvent structure. M Fujinaga, LTJ Delbaere, GD Brayer & MNG James, *L Mol Biol* **183**, 479–502 (1985).
- Halabi et al 2009:** Protein sectors: Evolutionary units of three-dimensional structure. N Halabi, O Rivoire, S Leibler & R Ranganathan, *Cell* **138**, 774–786 (2009).
- Mora & Bialek 2011:** Are biological systems poised at criticality? T Mora & W Bialek. *J Stat Phys* **144**, 268–302 (2011); arXiv:1012.2242 [q-bio.QM] (2010).
- Mora et al 2010:** Maximum entropy models for antibody diversity. T Mora, AM Walczak, W Bialek & CG Callan, *Proc Nat'l Acad Sci (USA)* **107**, 5405–5410 (2010).
- Russ et al 2005:** Natural-like function in artificial WW domains. WP Russ, DM Lowery, P Mishra, MB Yaffe & R Ranganathan, *Nature* **437**, 579–583 (2005).
- Socolich et al 2005:** Evolutionary information for specifying a protein fold. M Socolich, SW Lockless, WP Russ, H Lee, KH Gardner & R Ranganathan, *Nature* **437**, 512–518 (2005).
- Weigt et al 2009:** Identification of direct residue contacts in protein-protein interaction by message passing. M Weigt, RA White, H Szurmant, JA Hoch & T Hwa, *Proc Nat'l Acad Sci (USA)* **106**, 67–72 (2009).
- [Should really give some pointers to the problem of sequence alignment!](#)
[\[Check this against discussion and references in relevant part of Chapter Two!\]](#) The modern picture of transcriptional regulation traces its origins to Jacob & Monod (1961), another of the great and classic papers that still are rewarding to read decades after they were published. Their views were motivated primarily by studies of the *lac* operon, and the origins of these reach back to Monod's thesis (1942), which was concerned the phenomenology of bacterial growth. As recounted in Judson (1979), for example, the idea that genes turn on because of the release from repression was due to Szilard; the written record of these ideas is not as clear as it could be, but one can try Szilard (1960). For a modern view, faithful to the history, see Müller-Hill (1996). The other “simple,” paradigmatic example of protein-DNA interactions in the regulation of gene expression is the case of bacteriophage λ , which is reviewed by Ptashne (1986), which has also evolved with time (Ptashne 1992); see also Ptashne (2001). These systems provided the background for the pioneering discussion of sequence specificity in protein-DNA interactions (von Hippel & Berg 1986, Berg & von Hippel 1987, 1988). In parallel to this statistical approach, there were direct biochemical measurements of binding energies, and an early attempt to bring these different literatures into correspondence was by Stormo & Fields (1998).

- Berg & von Hippel 1987:** Selection of DNA binding sites by regulatory proteins. I: Statistical-mechanical theory and application to operators and promoters. OG Berg & PH von Hippel, *J Mol Biol* **193**, 723–743 (1987).
- Berg & von Hippel 1988:** Selection of DNA binding sites by regulatory proteins. II: The binding specificity of cyclic AMP receptor protein to recognition sites. OG Berg & PH von Hippel, *J Mol Biol* **200**, 709–723 (1988).
- von Hippel & Berg 1986:** On the specificity of DNA–protein interactions. PH von Hippel & OG Berg, *Proc Nat'l Acad Sci (USA)* **83**, 1608–1612 (1986).
- Jacob & Monod 1961:** Genetic regulatory mechanisms in the synthesis of proteins. F Jacob & J Monod, *J Mol Biol* **3**, 318–356 (1961).
- Judson 1979:** *The Eighth Day of Creation* HF Judson (Simon and Schuster, New York, 1979).
- Monod 1942:** *Recherche sur la Croissance des Cultures Bactériennes* J Monod (Hermann, Paris, 1942).
- Müller–Hill 1996:** *The lac Operon: A Short History of a Genetic Paradigm* B Müller–Hill (W de Gruyter & Co, Berlin, 1996).
- Ptashne 1986:** *A Genetic Switch: Gene Control and Phage λ* . M Ptashne (Cell Press, Cambridge MA, 1986).
- Ptashne 1992:** *A Genetic Switch, Second Edition: Phage λ and Higher Organisms*. M Ptashne (Cell Press, Cambridge MA, 1992).
- Ptashne 2001:** *Genes and Signals*. M Ptashne (Cold Spring Harbor Laboratory Press, New York, 2001).
- Stormo & Fields 1998:** Specificity, free energy and information content in protein–DNA interactions. GD Stormo & DS Fields, *Trends Biochem Sci* **23**, 109–113 (1998).
- Szilard 1960:** The control of the formation of specific proteins in bacteria and in animal cells. L Szilard, *Proc Nat'l Acad Sci (USA)* **46**, 277–292 (1960).
- The emergence of whole genome sequences opened several new approaches to the problem of specificity. One important idea is that sequences that are targets for protein binding should have a non-random structure, and we should be able to find this in a relatively unsupervised fashion (Bussemaker et al 2000a,b). [\[Need more here!\]](#)
- Bussemaker et al 2000a:** Building a dictionary for genomes: Identification of presumptive regulatory sites by statistical analysis. H Bussemaker, H Li & ED Siggia, *Proc Nat'l Acad Sci (USA)* **97**, 10096–10100 (2000).
- Bussemaker et al 2000b:** Regulatory element detection using a probabilistic segmentation algorithm. H Bussemaker, H Li & ED Siggia, *Proc Int Conf Intell Sys Mol Biol*—bf 8, 67–74 (2000).
- Need pointers to different large scale experimental approaches—protein binding arrays (Mukherjee et al 2004), ChiP, etc.. Circle back to work from Quake group (Maerkl & Quake 2007).** For an approach to the analysis of such measurements making explicit use of dimensionality reduction methods (Appendix **), see Kinney et al (2007). This approach inspired experiments aimed at wider exploration of sequence space (Kinney et al 2010). For other such explorations, see Ligr et al (2006) and Gertz et al (2009).
- Gertz et al 2009:** Analysis of combinatorial cis–regulation in synthetic and genomic promoters. J Gertz, ED Siggia & BA Cohen, *Nature* **457**, 215–218 (2009).
- Kinney et al 2007:** Precise physical models of protein–DNA interaction from high-throughput data. JB Kinney, G Tkačik & CG Callan Jr, *Proc Nat'l Acad Sci (USA)* **104**, 501–506 (2007).
- Kinney et al 2010:** Using deep sequencing to characterize the biophysical mechanism of a transcriptional regulatory sequence. JB Kinney, A Murugan, CG Callan Jr & EC Cox, *Proc Nat'l Acad Sci (USA)* **107**, 9158–9163 (2010).
- Ligr et al 2006:** Gene expression from random libraries of yeast promoters. M Ligr, R Siddharthan, FR Cross & ED Siggia, *Genetics* **172**, 2113–2122 (2006).
- Maerkl & Quake 2007:** A systems approach to measuring the binding energy landscape of transcription factors. SJ Maerkl & SR Quake, *Science* **315**, 233–237 (2007).
- Mukherjee et al 2004:** Rapid analysis of the DNA binding specificities of transcription factors with DNA microarrays. S Mukherjee, MF Berger, G Jona, XS Wang, D Muzzey, M Snyder, RA Young & ML Bulyk, *Nature Genetics* **36**, 1331–1339 (2004).
- Need to segue to discussions of evolvability etc.. Probably more references to cite!**
- Maerkl & Quake 2009:** Experimental determination of the evolvability of a transcription factor. SJ Maerkl & SR Quake, *Proc Nat'l Acad Sci (USA)* **106**, 18650–18655 (2006).
- Mustonen et al 2008:** Energy–dependent fitness: A quantitative model for the evolution of yeast transcription factor binding sites. V Mustonen, J Kinney, CG Callan Jr & M Lässig, *Proc Nat'l Acad Sci (USA)* **105**, 12376–12381 (2008).
- Sengupta et al 2002:** Specificity and robustness in transcription control networks. A Sengupta, M Djordjevic & BI Shraiman, *Proc Nat'l Acad Sci (USA)* , **99**, 2072–2077, (2002).

B. Ion channels and neuronal dynamics

The functional behavior of neurons involves the generation and processing of electrical signals. The dynamics of these currents and voltages are determined by the ion channels which sit in the cell membrane. As noted in our discussion of the rod photoreceptor cell (Section I.C), the cell membrane itself is insulating, and hence there would be no interesting electrical dynamics if not for specific conducting pores. These pores are protein molecules that can change their structure in response to various signals, including the voltage across the membrane, and this means that the system of channels interacting with the voltage constitutes a potentially complex nonlinear dynamical system. We can also think of the ion channels in the cell membrane as a network of interacting protein molecules, with the interactions mediated through the transmembrane voltage. In contrast to many other such biochemical systems, we actually know the equations that describe the network dynamics, and as a result the questions of fine tuning vs. robustness can be posed rather sharply.

When we move from thinking about individual neurons to thinking about circuits and networks of neurons,




FIG. 98 Activation curve for an ion channel, from Eq (505), with $Q = 4$.

which really do the business of the brain, it is easy to imagine that the neurons are ‘circuit elements’ with some fixed properties. We enhance this tendency by drawing circuit diagrams in which we keep track of whether neurons excite or inhibit one another, but nothing else about their dynamics is made explicit. In fact, our genome encodes $\sim 10^2$ different kinds of channels, each with its own kinetics, and this range is expanded even further by the fact that many of these channels have multiple subunits, and it is possible to splice together the subunits in different combinations. On the one hand, this creates enormous flexibility, and presumably adds to the computational power of the nervous system. On the other hand, this range of possibilities raises a problem of control. A typical neuron might have eight or nine different kinds of channels, and we will see that the dynamics of the cell depend rather sensitively on how many of each kind of channel is present. In keeping with the theme of this Chapter, it might seem that cells need to tune their channel content very precisely, yet this needs to happen in a robust fashion.

To explore the tradeoff between fine tuning and robustness in neurons, we need to understand the dynamics of the channels themselves. For simplicity, let’s neglect the spatial structure of the cell and assume we can talk about a single voltage difference V between inside and outside. Then since the membrane acts as a capacitor, we can write, quite generally,

$$C \frac{dV}{dt} = I_{\text{channels}} + I_{\text{ext}}, \quad (501)$$

where I_{ext} is any external current that is being injected (perhaps by us as experimenters) and I_{channels} is the current flowing through the channels. Each channel acts more or less as an Ohmic conductance, and the structure of the channel endows it with specificity for particular

ions. Since the cell works to keep the concentrations of ions different on the inside and outside of the cell, the thermodynamic driving force for the flow of current includes both the electrical voltage and a difference in chemical potential; it is conventional to summarize this by the ‘reversal potential’ V_i for the currents flowing through channels of type i , which might involve a mix of ions. Since current only flows through open channels, we can write

$$I_{\text{channels}} = - \sum_i g_i N_i f_i (V - V_i), \quad (502)$$

where g_i is the conductance of one open channel of type i , N_i is the total number of these channels, f_i is the fraction which are open, and V_i is the reversal potential. If each channel has just two states, open and closed, then their dynamics would be described by

$$\frac{df_i}{dt} = - \frac{1}{\tau_i(V)} [f_i - f_i^{\text{eq}}(V)]. \quad (503)$$

The equilibrium fraction of open channels as a function of voltage, $f_i^{\text{eq}}(V)$, often is called the activation curve, and $\tau_i(V)$ is the time constant for relaxation to this equilibrium.

What is a reasonable shape for the activation curve? We are describing a protein molecule that can exist in two states, and the equilibrium between these two states depends on voltage. This is possible only if the transition from closed to open rearranges the charges in the protein. In the simplest model, then, the opening of the channel effectively moves a charge Q across the membrane, and so the free energy difference between open and closed states will be $\Delta F = F_0 - QeV$. Then the equilibrium probability of a channel being open will be given by

$$f^{\text{eq}}(V) = \frac{1}{1 + \exp[(F_0 - QeV)/k_B T]} \quad (504)$$

$$= \frac{1}{1 + \exp[-(V - V_{1/2})/V_w]}, \quad (505)$$

where the point of half maximal activation is $V_{1/2} = F_0/(Qe)$, and the width of the activation curve is $V_w = k_B T/Qe$, as shown in Fig 98. The charge Q is referred to as the ‘gating charge.’ We recall that, at room temperature, $k_B T/e = 25$ mV, so that even with relatively small values of Q we expect channels to make the transition from closed to open in a window of ~ 10 mV or so. The location of the midpoint $V_{1/2}$ depends on essentially all aspects of the protein structure in the open and closed states, so it is harder to get intuition for this parameter. In practice, different channels have $V_{1/2}$ values in the range [look this up to give a meaningful survey ..].

It’s useful to think about the linearized dynamics; we imagine that there is some steady state at a ‘resting potential’ $V = V_0$, and study small perturbations around

this steady state. The full dynamics are

$$C \frac{dV}{dt} = - \sum_i g_i N_i f_i (V - V_i) + I_{\text{ext}}, \quad (506)$$

$$\frac{df_i}{dt} = - \frac{1}{\tau_i(V)} [f_i - f_i^{\text{eq}}(V)], \quad (507)$$

and the linearization is

$$C \frac{d\delta V}{dt} = - \sum_i g_i N_i f_i^{\text{eq}}(V) \delta V - \sum_i g_i N_i (V_0 - V_i) \delta f_i + I_{\text{ext}}, \quad (508)$$

$$\frac{d\delta f_i}{dt} = - \frac{1}{\tau_i(V_0)} \left[\delta f_i - \left. \frac{df_i^{\text{eq}}(V)}{dV} \right|_{V=V_0} \delta V \right]. \quad (509)$$

Fourier transforming, we can solve for the channel dynamics,

$$\frac{d\delta f_i}{dt} = - \frac{1}{\tau_i(V_0)} \left[\delta f_i - \left. \frac{df_i^{\text{eq}}(V)}{dV} \right|_{V=V_0} \delta V \right] \quad (510)$$

$$-i\omega \delta \tilde{f}_i(\omega) = - \frac{1}{\tau_i(V_0)} \left[\delta \tilde{f}_i(\omega) - \left. \frac{df_i^{\text{eq}}(V)}{dV} \right|_{V=V_0} \delta \tilde{V}(\omega) \right] \quad (511)$$

$$\delta \tilde{f}_i(\omega) = \frac{[df_i^{\text{eq}}(V)/dV]_0}{-i\omega + 1/\tau_i(V_0)} \delta \tilde{V}(\omega), \quad (512)$$

and then substitute,

$$C \frac{d\delta V}{dt} = - \sum_i g_i N_i f_i^{\text{eq}}(V) \delta V - \sum_i g_i N_i (V_0 - V_i) \delta f_i + I_{\text{ext}}$$

$$-i\omega C \delta \tilde{V}(\omega) = - \sum_i g_i N_i f_i^{\text{eq}}(V) \delta \tilde{V}(\omega) - \sum_i g_i N_i (V_0 - V_i) \delta \tilde{f}_i(\omega) + \tilde{I}_{\text{ext}}(\omega) \quad (513)$$

$$-i\omega C \delta \tilde{V}(\omega) = - \sum_i g_i N_i f_i^{\text{eq}}(V) \delta \tilde{V}(\omega) - \sum_i \frac{[g_i N_i (V_0 - V_i) df_i^{\text{eq}}(V)/dV]_0}{-i\omega + 1/\tau_i(V_0)} \delta \tilde{V}(\omega) + \tilde{I}_{\text{ext}}(\omega). \quad (514)$$

Collecting terms, we find

$$\left[-i\omega C + \frac{1}{R_0} + \sum_i \frac{g_i N_i (V_0 - V_i) [df_i^{\text{eq}}(V)/dV]_0}{-i\omega + 1/\tau_i(V_0)} \right] \delta \tilde{V}(\omega) = \tilde{I}_{\text{ext}}(\omega). \quad (515)$$

The resting resistance of the membrane is defined by

$$\frac{1}{R_0} = \sum_i g_i N_i f_i^{\text{eq}}(V). \quad (516)$$

The term in brackets in Eq (515) is the inverse impedance (or “admittance”) of the system.

To understand what is going on here, it’s useful to think about channels which have fast ($1/\tau_i \gg \omega$) or slow ($1/\tau_i \ll \omega$) responses. The fast channels renormalize the resistance,

$$\frac{1}{R_0} \rightarrow \frac{1}{R_0} + \sum_{i \in \text{fast}} \tau_i(V_0) g_i N_i (V_0 - V_i) \left. \frac{df_i^{\text{eq}}(V)}{dV} \right|_{V=V_0}. \quad (517)$$

Importantly, the correction to the resistance can be either positive or negative. Suppose that, as in Fig 98, the channels tend to open in response to increasing voltage, as most channels do. Then $[df_i^{\text{eq}}(V)/dV]_0 > 0$. But if this channel is specific for an ion with a reversal potential above the resting potential ($V_i > V_0$), then opening the channel creates a stronger tendency to pull the voltage toward this higher potential, which is a regenerative effect—a negative resistance.

If the channels are slow, they make a contribution to the imaginary part of the admittance, along with the

capacitance,

$$-i\omega C \rightarrow -i\omega C + \frac{1}{-i\omega} \sum_{i \in \text{slow}} g_i N_i (V_0 - V_i) \left. \frac{df_i^{\text{eq}}(V)}{dV} \right|_{V=V_0}. \quad (518)$$

Again the sign depends on details. If the channels are opened by increasing voltage and the reversal potential is *below* the resting potential, then their contribution is (almost) like an inductance, and can generate a resonance by competing with the capacitance. This resonance is at a frequency

$$\omega_* = \left[\frac{1}{C} \sum_{i \in \text{slow}} g_i N_i (V_0 - V_i) \left. \frac{df_i^{\text{eq}}(V)}{dV} \right|_{V=V_0} \right]^{1/2} \quad (519)$$

which, interestingly, does not depend on the precise value of the time constants defining the channel kinetics, although one must obey the condition $\omega_* \gg 1/\tau_i(V_0)$ for all $i \in \text{slow}$.

Problem 99: Equivalent circuits. Equation (515) shows that each type of channel contributes a parallel path for current flow through the membrane. The impedance of this path is defined by

$$\frac{1}{\tilde{Z}_i(\omega)} = g_i N_i f_i^{\text{eq}}(V) + \frac{g_i N_i (V_0 - V_i) [df_i^{\text{eq}}(V)/dV]_0}{-i\omega + 1/\tau_i(V_0)}. \quad (520)$$

Without resorting to the fast/slow approximations above, draw an equivalent circuit using the standard lumped elements (capacitance, resistance, inductance) which realizes this impedance. Show how the parameters of the lumped elements relate to the parameters of the channels.

So, we have seen that even in response to small signals, the dynamics of ion channels generate an interesting complement of electronic parts: resistors, inductors, and negative resistors. Certainly one can put these together to make a filter, playing the effective inductance of the channels against the intrinsic capacitance of the membrane, as noted above. With the negative resistor one can sharpen the resonance, and even generate an instability; presumably on the other side of the instability is a genuine oscillator.

Problem 100: Oscillations. Construct a minimal model for ion channels in the cell membrane that supports a stable, limit cycle oscillation of the voltage.

The negative resistance alone means that we can have (without oscillations) an instability of the steady state around which we were expanding, presumably because the real system is multi-stable. To see this more clearly, consider just two types of channels—a ‘leak’ channel which is open independent of the voltage and has a reversal potential of zero, and some other channel which opens in response to increasing voltage. Then the dynamics are

$$C \frac{dV}{dt} = -G_{\text{leak}} V - gN f(V - V_r), \quad (521)$$

$$\frac{df}{dt} = -\frac{1}{\tau(V)} [f - f_{\text{eq}}(V)]. \quad (522)$$

The steady state solutions are determined by solving two simultaneous equations, usually called the nullclines, obtained by setting the time derivatives equal to zero:

$$f = f_{\text{eq}}(V) \quad (523)$$

$$V = V_r \frac{f}{f + G_{\text{leak}}/gN}; \quad (524)$$

these are shown schematically in Fig 99, for some reasonable choice of parameters. Evidently there are three solutions to the two simultaneous equations, and it is fairly easy to show that two are stable and one is unstable. The two stable states correspond, roughly, to one state in which all the channels are closed and the voltage is zero (the reversal potential of the leak), and one state in which all the channels are open and the voltage is near the reversal potential for these channels. The bistability means that, if the cell starts in the low voltage state, injection of a relatively small, brief current can drive the system across a threshold (separatrix) so that it falls into the high voltage state after the current pulse is complete. This is a form of memory (interesting, although not very realistic), but also a substantial amplification of the incoming signal, especially if the parameters are tuned so that the difference in voltage to the unstable state is small.

Problem 101: Bistability. Work through a concrete example of the ideas in the previous paragraphs, perhaps using the detailed model from Fig 99. You should be able to verify, analytically, the claims about stability of the different steady states. Explain how these analytic criteria can be converted into a test for stability of each steady state that can be ‘read off’ directly from the plots in Fig 99. Analyze the response to brief pulses of current, showing that there is a well defined threshold for switching from one stable state to the other.

All the different kinds of dynamics we have seen thus far—filtering, oscillation, and bistability—can be generated by just one kind of channel with only two states.




FIG. 99 Bistability in a simple model of a neuron. The channel nullcline is Eq (523), and the voltage nullcline is Eq (524). To be explicit we choose $f_{\text{eq}}(V)$ from Eq (505), with $V_{1/2} = 70$ and $V_w = 10$, and $G_{\text{leak}}/gN = 0.1$. Note that there are three crossing points, corresponding to steady states. The low voltage and high voltage states are stable; the intermediate voltage state is unstable.

Real neurons are much more complex. One important class of dynamics that we can't quite see in the simplest models is 'excitability.' In this case, a small pulse again drives the system across a threshold, but what would have been a second stable state is destabilized by relaxation of some other degrees of freedom; the result is that the system takes a long, and often stereotyped, trajectory through its phase space before coming back to its original steady state after the input pulse is over. The action potential is an example of such excitable dynamics [should we have a sketch of what this means in a simple phase plane?].

Our understanding of ion channels goes back to the classic work of Hodgkin and Huxley in the 1940s and 50s. They studied the giant axon, a single cell, visible to the naked eye, which runs along the length of a squid's body, and along which action potentials are propagated to trigger the squid's escape reflex. Passing a conducting wire through the interior of the long axon, they short-circuited the propagation, insuring that the voltage across the membrane was spatially uniform, as in our idealization above. They then studied the current that flowed in response to steps of voltage. If the picture of channels is correct, then with the voltage held constant, there should be an (Ohmic) flow of current through the open channels. If we step suddenly to a new value of the voltage, Ohm's law tell us that the current through the open channels will change immediately, but there will be a prolonged time dependence that results from the open or closing of channels as they equilibrate at the new voltage. In the simple model with two states, this changing

current should relax exponentially to a new steady state; in particular, the initial slope of the current should be finite.

Hodgkin and Huxley found that the relaxation of the current at constant voltage has a gradual start, as if the channels had not one closed state but several, and the molecules had to go through these states in sequence before opening. They chose to describe these dynamics of the currents by imagining that, in order for the channel to be open, there were several independent molecular "gates" that all had to be open. Each gate could have only two states, and would obey simple first order kinetics, but the probability that the channel is open would be the product of the probabilities that the gates were open. In the simple case that the multiple gates are identical, the probability of the channel being open is just a power of the 'gating variable' describing the probability that one gate is open. Hodgkin and Huxley also discovered that at least one important class of channels open in response to increased voltage, and then seem to close over time. They described this by saying that in addition to 'activation gates' that were opened by increasing voltage, there were 'inactivation gates' which closed in response to increasing voltage, but these had slower kinetics. Putting the pieces together, they described the fraction of open channels as

$$f_i = m_i^{\alpha_i} h_i^{\beta_i}, \quad (525)$$

where m and h are activation and inactivation gates, respectively, and the powers α and β count the number of these gates that contribute to the opening of one channel. The kinetics are then described by

$$\frac{dm_i}{dt} = -\frac{1}{\tau_i^{(m)}(V)} [m_i - m_i^{\text{eq}}(V)] \quad (526)$$

$$\frac{dh_i}{dt} = -\frac{1}{\tau_i^{(h)}(V)} [h_i - h_i^{\text{eq}}(V)], \quad (527)$$

and finally the voltage (again neglecting spatial variations) obeys

$$C \frac{dV}{dt} = -\sum_i g_i N_i m_i^{\alpha_i} h_i^{\beta_i} (V - V_i). \quad (528)$$

Problem 102: Two gates. Suppose that each channel has two independent structural elements ("gates"), each of which has two states. Assuming that the two gates are independent of one another, fill in the steps showing that the dynamics of the channels are as described above. In particular, show that after a sudden change in voltage, the fraction of open channels starts to change as $\propto t^2$, not $\propto t$ as expected if the entire channel only has two states. [This, and the preceding paragraph, might be a little too telegraphic. Need feedback here!]

Problem 103: Hodgkin and Huxley revisited. The original equations written by Hodgkin and Huxley are as follows:⁶⁹

$$C \frac{dV}{dt} = -\bar{g}_L(V - V_L) - \bar{g}_{\text{Na}}m^3h(V - V_{\text{Na}}) - \bar{g}_K n^4(V - V_K) + I(t) \quad (529)$$

$$\frac{dn}{dt} = (0.01V + 0.1) \exp(-V/10)(1 - n) - 0.125n \exp(V/80)n \quad (530)$$

$$\frac{dm}{dt} = (0.1V + 2.5) \exp(-V/10 - 1.5)(1 - m) - 4 \exp(V/18)m \quad (531)$$

$$\frac{dh}{dt} = 0.07 \exp(V/20)(1 - h) - \exp(-V/10 - 4)h, \quad (532)$$

where Na and K refer to sodium and potassium channels, respectively; time is measured in milliseconds and V is measured in millivolts. These equations are intended to describe a small patch of the membrane, and so many parameters are given per unit area: $C = 1 \mu\text{F}/\text{cm}^2$, $\bar{g}_L = 0.3 \text{mS}/\text{cm}^2$, $\bar{g}_{\text{Na}} = 120 \text{mS}/\text{cm}^2$, and $\bar{g}_K = 36 \text{mS}/\text{cm}^2$; the reversal potentials are $V_L = 10.613 \text{mV}$, $V_{\text{Na}} = 115 \text{mV}$, and $V_K = -12 \text{mV}$.

(a.) Rewrite these equations in terms of equilibrium values and relaxation times for the gating variables, e.g.

$$\frac{dm}{dt} = -\frac{1}{\tau_m(V)} [m - m_{\text{eq}}(V)]. \quad (533)$$

Plot these quantities. Can you explain, intuitively, the form of the curves?

(b.) Simulate the dynamics of the Hodgkin–Huxley equations in response to constant current inputs. Show that there is a threshold current, above which the system generates period pulses. Explore the frequency of the pulses as a function of current.

(c.) Suppose that the injected current consists of a mean (less than the threshold you identified in [b]), plus a small component at frequency ω . By some appropriate combination of analytic and numerical methods, find the impedance $Z(\omega)$ for different values of the mean injected current. Show that the membrane has a resonance, and explore what happens to this resonance as the mean current is increased toward threshold. How do your results connect to the frequency of pulses above threshold?

(d.) Real axons are essentially long thin cylinders. Show that, if we allow the voltage to vary along the length of the axon, there should be a current per unit area flowing across the membrane of

$$I = \frac{a}{2R} \frac{\partial^2 V}{\partial z^2}, \quad (534)$$

where z is the coordinate along the cylinder, a is its radius, and R is the resistivity of the fluid filling the axon, assuming that resistance outside the axon is negligible. For the squid giant axon, $a \sim 250 \mu\text{m}$ and $R \sim 35 \Omega \cdot \text{cm}$. Use this result to write equations for the voltage and gating variables along the axon. Note that only the dynamics of voltage is sensitive to spatial derivatives. Why?

(e.) Simulate the response of a long segment of the axon to a current pulse injected at one end. Show that small pulses result in spatially restricted voltage responses, while larger pulses produce a




FIG. 100 The action potential that emerges from the Hodgkin–Huxley model. **Need to decide what to say, what other things to reproduce ...**

propagating pulse. Confirm that these pulses become more stereotyped as the propagate, and have a velocity that is independent of the input current. What is this velocity? How does it compare to the observed speed of action potentials, $v \sim 20 \text{m/s}$?

Problem 104: Simplification. It is very hard to make analytic progress in understanding the dynamics of a system with five variables. There is a history of trying to approximate the system by exploiting the fact that the different variables have very different time scales. See how far you can go along this path. I have left this problem deliberately open-ended. For one approach, see Abbott and Kepler (1990).

It is good to pause here and review how we know that the Hodgkin–Huxley description of ion channels is correct. **[Not sure how much of this should be illustrated by figures from the original papers?]** The initial triumph, which you are asked to reproduce in the problem above, is the prediction of the propagating action potential itself, as in Fig 100, with the correct speed. The model also predicts that, as the action potential passes, there is a net flux of potassium and sodium across the membrane. On long time scales, this must be balanced by the action of pumps that maintain the concentration differences between the inside and outside of the cell. But either by looking quickly or by poisoning the pumps, one should be able to detect the flux, for example using radioactive tracers, and this works, quantitatively.

[This is all a little vague; should go back and try to do better!] Nature provides a variety of toxins which block the action potential in different ways, and we can also find artificial blockers, for example using ions with very large radius that can literally plug the hole in open channels. It is striking that these agents act selectively on different channels, and one can verify that this way of isolating the dynamics of sodium and potassium channels matches the Hodgkin–Huxley description. If we can arrange for the channels to “open” but be blocked, then the structural change of the channel molecule upon opening should still

⁶⁹ The only difference from the original paper is that we use the modern sign convention for the voltage. Notice that this original formulation is in terms of a “maximal conductance” for each type of “current,” while in modern language we could talk about the number of each type of channel. In fact, the more phenomenological description persists, because it corresponds more directly to what is measured, but this allows us to forget that parameters such as \bar{g}_K actually measure the number of copies of a protein that have been inserted into the membrane.

move the gating charge across the membrane, and if we are careful this should be measurable essentially as a delayed capacitive response to changes in the applied voltage. These “gating currents” have indeed been detected, and in some cases it has been possible to match these quantitatively not only against predictions based on the form of the activation curve, but also to genetically engineer the channels and show that changes in the activation curve and gating currents track one another. [How much detail here? Give the example of shaker?]

If individual channels are independent of one another, then their opening and closing events should be independent. If we look at a small patch of the membrane, there will not be that many channels present, and we might be able to see that the discrete events in the individual molecules don't quite average out—there should be noise from the random opening and closing of the single channels. This channel noise has been detected, and has the spectral properties expected from the Hodgkin–Huxley model. Finally, if we look at even smaller patches of the membrane, and have proportionately more sensitive amplifiers, we should be able to see the opening and closing of single channels. Again, this works. Most importantly, we can look at the distribution of times that individual channels spend in the open and closed states, and connect this to the kinetics predicted by the Hodgkin–Huxley model and its generalizations. Although these more detailed measurements have revealed new features of channel kinetics even in well studied examples, in outline the picture given to us by Hodgkin and Huxley has stood the test of time. [Again, should probably show some figures. Emphasize how remarkable it is to be looking at individual molecular events—current flow through sub-nanometer pores! Maybe even discuss shot noise through open channels?]

Problem 105: Channel noise. Give a problem that maps the HH model onto a stochastic picture of channel states, and then derive the expected properties of the channel noise. Remember that we did the simplest version of this in Chapter 1.

Problem 106: Single channel kinetics. Give a problem that explores how single channel kinetics are connected to the macroscopic kinetics.

Now that we have confidence in our mathematical description of neurons, it is time to realize now just how many parameters are involved. A typical cell expresses eight or nine different kinds of channels. Each channel is described by the dynamics of two gating variables. If we imagine that activation or inactivation curves have the simple sigmoidal form as in Fig 99, then there are roughly two parameters for each such curve—the voltage

at half activation and the slope or width—and at least one more parameter to set the time scale of the kinetics. Finally, there is the total number of channels, or the maximum conductance achieved if all the channels are open. All together, then, this is ~ 7 parameters per channel type, or roughly fifty parameters for the entire neuron, conservatively. Importantly, to a large extent the cell actually has control over these parameters, and, in a meaningful sense, can adjust them almost continuously.

How do these adjustments occur? Most obviously, the total number of open channels is controlled in the same that all other protein copy numbers are controlled. Sometimes, because of the clearer connection to experiment, one speaks about the ‘maximal conductance’ associated with a particular type of channel ($G_i^{\max} = g_i N_i$), but this obscures the fact that this parameter really is the total number of copies of the protein that the cell has expressed and inserted into the membrane. The parameters of the activation curves and the time constants are intrinsic properties of the proteins, but these too can be adjusted in several ways. First, like all proteins, ion channels can be covalently modified by phosphorylation etc.. More importantly, the genome encodes a huge number of different ion channels proteins; the human genome has 90 different potassium channels alone. While these do form classes based on their dynamics, there is considerable variation within classes, and since many of these genes have multiple alternative splicings, there is the potential for almost continuous parameter variation. These different mechanisms of variation interact; as an example, different splicing variants can exhibit different sensitivity to phosphorylation.

Problem 107: Continuous adjustment of electrical dynamics. [It might be that I should take the students by hand through the model; let's see how this works.] To illustrate the possibility of nearly continuous adjustments in the electrical dynamics of neurons, consider the case of the hair cells in the turtle ear. In these cells (cf Section 2.5), one contribution to frequency selectivity comes from a resonance in the electrical response of the hair cell itself. This resonance is driven by a combination of voltage-gated calcium channels and calcium-activated potassium channels. There is a detailed model of this system, described by Wu & Fettplice (2001). Try to understand what they have done, and reproduce the essential theoretical results. In particular, what is the role of “details” (e.g., the building of channels out of combinations of different subunits) in generating the correct qualitative behavior?

One well studied example of channel dynamics is in the stomatogastric ganglion of crabs and lobsters, schematized in Fig 101. This is a network of ~ 30 neurons which generates a rhythm, and this rhythm in turn drives muscles which actuate teeth in the crab stomach, grinding its




FIG. 101 The stomatogastric ganglion (STG) in crustaceans, from Marder & Bucher (2007). At top left, the location of the STG and the commissural ganglion (CoG) in a lobster. At top right, a schematic of the ganglion dissected out of the animal, and the opportunities for recording the activity of the neurons. At bottom, simultaneous extracellular recordings from nine motor nerves at the output of this network. Names indicate particular neurons which can be identified in each individual (as with the named neurons in the fly visual system discussed in [pointer]), and in some cases (e.g., avn, mvn) we can identify spikes from several individual neurons in the recording from one nerve. There are two main rhythms, the faster pyloric rhythm in cells PD, LP, PY, VD and IC, and the slower gastric mill rhythm in cells MG, DG, GM, LPG and LG.

food. Evidently getting the correct rhythm is important in the life of the organism. If one records the electrical signals from individual neurons, several of the cells produce period bursts of action potentials, and a handful of cells are ‘pacemakers’ that can generate this periodic pattern without input from the other cells. In one such cell (the lateral pyloric neuron), experiments show that there are seven different channel types. An important feature of this cell, shared by many other cells, is the presence of voltage-gated calcium channels. This means that, as action potentials occur, they trigger calcium flux into the cell. Because there are also channels which are directly affected by the calcium concentration, a complete model must include a description of the calcium buffering or pumping that counterbalances this flux.

It is worth being very explicit about all these ingredients in the dynamics of the lateral pyloric neuron, not least to get a sense for the state of the art in such analyses. As before, we will neglect the spatial structure of the cell, so there is just one relevant voltage difference V between the inside and outside of the cell, which obeys a slight generalization of Eq (528),

$$C \frac{dV}{dt} = - \sum_i g_i N_i m_i^{\alpha_i} h_i^{\beta_i} (V - E_i) + I_{\text{ext}}, \quad (535)$$

where I_{ext} is any externally injected current and E_i is the reversal potential for channel type i . The kinetics of the gating variables m_i and h_i are governed by Eq’s (526) and (527), respectively. For most of the channels, we can take the equilibrium values of the gating variables to be

channel type	$g_i N_i$ (μS)	E_i (mV)	midpoints (mV)	widths (mV)	rates (s^{-1})
i = 1: “delayed rectifier”	0.35	$E_K = -80$			
activation equilibrium ($\alpha_1 = 4$)			$V_{1/2}^{m_1} = -25$	$V_w^{m_1} = 17$	
activation kinetics			$V_1^{(m)} = 10$	$1/\gamma_1^{(m)} = 22$	$k_1^{(m)} = 180$
i = 2: Ca^{++} current 1	0.21	E_{Ca}			
activation ($\alpha_2 = 1$)			$V_{1/2}^{m_2} = -11$	$V_w^{m_2} = 7$	50
inactivation ($\beta_2 = 1$)			$V_{1/2}^{h_2} = -50$	$V_w^{h_2} = -8$	16
i = 3: Ca^{++} current 2	0.047	E_{Ca}			
activation ($\alpha_3 = 1$)			$V_{1/2}^{m_3} = -22$	$V_w^{m_3} = 7$	10
i = 4: “inward rectifier”	0.037	-10			
activation equilibrium ($\alpha_4 = 1$)			$V_{1/2}^{m_4} = -70$	$V_w^{m_4} = -7$	
activation kinetics			$V_4^{(m)} = -110$	$1/\gamma_1^{(m)} = 13$	$k_1^{(m)} = 0.33$
i = 5: “leak”	0.1	-50			
i = 6: “A-current”	2.2	$E_K = -80$			
activation equilibrium ($\alpha_6 = 3$)			$V_{1/2}^{m_6} = -12$	$V_w^{m_6} = 26$	
activation kinetics					$k_6^{(m)} = 140$
inactivation equilibrium ($\beta_{6a} = 1$)			$V_{1/2}^{h_{6a}} = \dots$	$V_w^{m_{6a}} = \dots$	
inactivation kinetics					$k_{6a}^{(h)} = \dots$
inactivation equilibrium ($\beta_{6b} = 1$)			$V_{1/2}^{h_{6b}} = \dots$	$V_w^{m_{6b}} = \dots$	
inactivation kinetics					$k_{6b}^{(h)} = \dots$

TABLE I A subset of channels in the lateral pyloric neuron, from Buchholtz et al (1992). For the delayed rectifier and the second type of calcium channel, there is no evidence for inactivation. The negative value of $V_w^{(h_2)}$ means, from Eq (505), that the probability of the inactivation gate being “open” decreases with increasing voltage. For calcium channels, the reversal potential varies, depending on the calcium concentration inside the cell, as in Eq (541), and the relaxation times do not have a detectable voltage dependence. The voltage dependence of the inward rectifier kinetics is opposite to Eq (538), that is $1/\tau \propto 1 + \exp[-\gamma_i^{(m)}(V - V_i^{(m)})]$. The leak current, by convention, is the current that exhibits no voltage or time dependence of its conductance. [Get details of the A-current right!](#)

given by the generalization of Eq (505),

$$m_i^{\text{eq}}(V) = \frac{1}{1 + \exp[-(V - V_{1/2}^{m_i})/V_w^{m_i}]}, \quad (536)$$

$$h_i^{\text{eq}}(V) = \frac{1}{1 + \exp[-(V - V_{1/2}^{h_i})/V_w^{h_i}]}, \quad (537)$$

and the time constants for relaxation of the gating variables are, phenomenologically,

$$\frac{1}{\tau_i^{(m)}(V)} = \frac{k_i^{(m)}}{1 + \exp[-\gamma_i^{(m)}(V - V_i^{(m)})]}, \quad (538)$$

$$\frac{1}{\tau_i^{(h)}(V)} = \frac{k_i^{(h)}}{1 + \exp[-\gamma_i^{(h)}(V - V_i^{(h)})]}. \quad (539)$$

As shown in Table I, this description works for several channel types, one selective for potassium, two for calcium, and one mixed, plus a “leak” that exhibits no significant time or voltage dependence of its conductance.

Two of the important channel types allow calcium to flow into the cell. As we will see, this current is big enough to change the concentration of calcium inside

the cell, and this has a variety of effects on other processes, including one of the channels that doesn’t fit the simple description we have given so far. So, we will need to describe the dynamics of the calcium concentration itself. The simplest model is that the calcium relaxes back to some internally determined steady state, $[\text{Ca}]_0 = 0.05 \mu\text{M}$, with a rate $k_{\text{Ca}} = 360 \text{s}^{-1}$, in which case

$$\frac{d[\text{Ca}]}{dt} = -k_{\text{Ca}} ([\text{Ca}] - [\text{Ca}]_0) + AI_{\text{Ca}}, \quad (540)$$

where I_{Ca} is the total calcium current ($I_{\text{Ca}} = I_2 + I_3$ from Table I). The constant $A = 300 \mu\text{M}/\text{nC}$ is inversely proportional to the volume into which the current flows, which experimentally comes out to be much smaller than the total volume of the cell body. As the concentration of calcium changes, the reversal potential for the calcium currents also changes,

$$E_{\text{Ca}} = \frac{k_B T}{2e} \ln \left(\frac{[\text{Ca}]_{\text{out}}}{[\text{Ca}]} \right), \quad (541)$$

where the calcium concentration outside the cell is $[\text{Ca}]_{\text{out}} = 13 \text{mM}$.

We are still missing three of the channel types in this cell. First, there is another potassium channel that is almost described by our standard model, but the inactivation seems to involve two processes that occur on different time scales. This can be captured by replacing

$$h_6 \rightarrow x(V)h_{6a} + [1 - x(V)]h_{6b}, \quad (542)$$

where the weighting function

$$x(V) = \frac{1}{1 + \exp[-(V - 7)/15]}, \quad (543)$$

$$m_7 = m_7^{\text{eq}}(V) = \frac{1}{1 + \frac{136}{V+6} (\exp[-(V + 34)/13] - \exp[-(V - 0.07)/7.9])}, \quad (544)$$

where V again is measured in mV [Need to check this carefully!]. The inactivation gates obey

$$\frac{dh_7}{dt} = a_7(V)(1 - h_7) - b_7(V)h_7, \quad (545)$$

where the rates

$$a_7(V) = 40 \exp[-(V + 39)/8], \text{ and} \quad (546)$$

$$b_7(V) = \frac{500}{1 + \exp[-(V + 40)/5]}, \quad (547)$$

are measured in s^{-1} . The total conductance that is contributed by these channels is large, $g_7 N_7 = 2300 \mu\text{S}$, al-

with V measured in mV as before.

Next, there is a fast sodium channel not unlike the ones that Hodgkin and Huxley found in the squid giant axon, with $\alpha_7 = 3$ and $\beta_7 = 1$. The activation is sufficiently fast that it can be approximated as instantaneous, so that m_7 is always at its equilibrium value, which varies with voltage in a slightly more complicated way than for the other channels,

though they are only open briefly.

The last type of channel, like the first two in Table I, is selective for potassium ions, but the probability of the channel being open is modulated by the intracellular calcium concentration. This channel has $\alpha_8 = \beta_8$, and the equilibrium state of the inactivation gate depends only on the calcium concentration,

$$h_8^{\text{eq}} = \frac{1}{1 + [\text{Ca}]/(0.6 \mu\text{M})}. \quad (548)$$

The equilibrium state of the activation gate, in contrast, depends both on voltage and on calcium,

$$m_8^{\text{eq}} = \frac{1}{1 + \exp[-(V + f[\text{Ca}])/23]} \cdot \frac{1}{1 + \exp[-(V + 16 + f[\text{Ca}])/5]} \cdot \frac{[\text{Ca}]}{2.5 \mu\text{M} + [\text{Ca}]}, \quad (549)$$

where $f = 0.6 \text{ mV}/\mu\text{M}$. The relaxation rates $k_8^{(m)} = 600 \text{ s}^{-1}$ and $k_8^{(h)} = 35 \text{ s}^{-1}$ show little if any voltage dependence. This seems like a complicated model, but it fits the experimental results very well, as in Fig 102.

Problem 108: Calcium dependent potassium conductances. Develop a microscopic picture to explain the combination of voltage and calcium dependences seen in Eq's (548) and (549). Remember that these equations describe the equilibrium fractions of molecules in particular states, so you need to relate these back to the free energies of the different states. Connect your discussion with the MWC models discussed in Appendix A.4 and [elsewhere?].

The model of the lateral pyloric neuron which we have described here represent the culmination of many years of effort, both in experiments on this particular system and in the exploration of these fully realistic generalizations of the Hodgkin–Huxley model to what seems the more typical case, with many different channel types functioning together. This model also represents a level of detail and complexity that I have tried to avoid so far, so some explanation is called for. First, the complexity consists largely of variations on a theme. Many channels are known to be described by the general picture of multiple activation and inactivation gates, so this provides a framework within which each new type of channel can be fit. Second, the complexity is justified by a large body of data. There are independent experiments on other systems, exploring quantitatively each of the types of channels that we see in this neuron, and detailed experiments

on this one cell to tease out the contributions of each of the channel types.

Problem 109: Justifying complexity. Go through Golowasch & Marder (1992), Buchholtz et al (1992), and Golowasch et al (1992), and explain the justification for each of the channel types in the model discussed above.

Indeed, the program of describing the electrical dynamics of single neurons in terms of generalized Hodgkin–Huxley models, usually with many different channel types functioning together, became a small industry. It really worked. In some cases one could go so far as to characterize the kinetics of particular channel types through measurements on single molecules, and then put these single molecule properties together to reproduce the functional behavior of the cell as a whole. This really is quite a beautiful body of work, and implements what many people would like to do in other systems, building from measured properties of individual molecular events up to macroscopic biological function. As emphasized above, we can think of the ion channels in the cell membrane as a network of interacting proteins, where the interaction is mediated by the voltage across the membrane rather than direct protein–protein encounters, and where the equations for the dynamics of the individual channels have a firm foundation. It is not unreasonable to claim that ion channels in the cell membrane are in fact the




FIG. 102 Dynamics of the calcium dependent potassium current, from Buchholtz et al (1992). Experimental data (noisy traces) from Golowasch & Marder (1992), solid lines from the model including Eq's (548) and (549). [Go back and understand how they isolate this contribution to the current]




FIG. 103 Simulations of a detailed model, with seven types of channel, for the lateral pyloric neuron in the stomatogastric ganglion of the crab. Changes in the pattern of activity as a function of the numbers of two different kinds of channel, where channel number here is expressed as the maximal conductance when all channels are open. Note that relatively small changes in these parameters can result in both quantitative and qualitative changes in the pattern of electrical activity, running the full range from silence to single spike firing to bursting. From Le Masson et al (1993).

best understood examples of biochemical networks, although the language typically used in describing these systems obscures this connection.

Despite their success, it came to be known, though not widely commented upon, that these models of coupled ion channel dynamics had a problem. While experiments often characterize the activation curves and kinetics of the individual channels, it is hard to make independent measurements of the total number of channels, or equivalently the maximum conductance when all the channels are open. Thus, one is left adjusting these parameters, trying to fit the overall electrical dynamics of the neuron—for example, the rhythmic bursting of the pyloric neuron. This fitting turns out to be delicate; as one adjusts the (many) parameters, one finds bifurcations to qualitatively different behaviors in response to relatively small changes. An example of this is shown in one two–dimensional slice through the seven dimensional space of channel numbers in the pyloric model, at the top in Fig 103.

Frankly, from a physicist's point of view this all seems a mess. There are many details one has to keep track of, and many parameters to adjust. One might be tempted just to walk away, and count this as a part of biology we don't want to know about. But there is a deep ques-

tion here:⁷⁰ if we have trouble adjusting the parameters of our models in order to reproduce the observed functional behaviors of particular cells, how do the cells themselves adjust these parameters to achieve their correct functions? How does it choose the ‘correct’ number of each type of channel to express? One could imagine that the cell has some sort of lookup table—I am a cell of type α , so I should express N_1^α molecules of channel type 1, N_{37}^α molecules of channel type 37, and so on. This is a bit implausible. More likely would be that the cell has a way of monitoring its activity and asking “how close am I to doing the right thing?,” generating an error signal that could be used to drive changes in the expression of the channels or perhaps their insertion into the membrane.

How can a neuron “know” whether it is exhibiting the desired pattern of electrical activity? It would need some signal that couples voltage changes across the membrane, which are quite fast, to the biochemical events regulating gene expression, which are quite slow. One idea is to use the intracellular calcium concentration as an intermediary. We know that many cellular processes are regulated by calcium, so one end of this is easy to imagine. But in the models described above the calcium concentration is an explicit part of the dynamics, so we can calculate, for example, the time average calcium concentration as function of the parameters of the model. What we see in Fig 104 is that $[Ca^{++}]$ does an excellent job of tracing the pattern of electrical activity in this cell. Thus if the system wants to stabilize a pattern of rhythmic bursting, it can do so via feedback mechanisms which try to hold the calcium concentration near a target value of $C_0 \sim 0.2 \mu\text{M}$.

Let us suppose that the expression of each channel protein is regulated by calcium, so that

$$\tau_i \frac{dN_i}{dt} = N_i^{\max} f_i([Ca^{++}]/C_0) - N_i, \quad (550)$$

where $f_i(x)$ is a sigmoidal function such as

$$f_i(x) = \frac{1}{1 + x^{\pm n}}. \quad (551)$$

Of course these equations have their steady state at $N_i = N_i^{\max} f_i([Ca^{++}]/C_0)$, but the calcium concentration must be determined self-consistently through the full dynamics of the channels and voltage. We should choose the signs of the calcium dependences to insure stability: channels which allow excitatory currents to flow will tend to drive increases in $[Ca^{++}]$, and so these should be




FIG. 104 Mean calcium concentration follows the pattern of electrical activity. Main figure shows the mean calcium concentration as a function of the same two variables shown in Fig 103. Small figure at right shows that the region of bursting activity corresponds almost perfectly to the region of parameter space in which the mean calcium concentration is between 0.1 and 0.3 μM , so that holding the calcium level fixed will stabilize bursting. From Le Masson et al (1993).

opposed by a decreasing function $f_i(x)$, and vice versa. Once we do this, if the regulation functions are steep [large value of n in Eq (551)], and the maximum possible numbers of channels (N_i^{\max}) are large, the dynamics will always be pulled into regimes where $[Ca^{++}] \approx C_0$. **We need a figure which illustrates this!**

Problem 110: A simple example of a self-tuning neuron. Need to find the simplest example of these models, and let the students work it through for themselves.

How can we tell if something like this sort of self-tuning really is happening? If neurons knew how many of each kind of channel to make, then they would try to do this no matter what the conditions. For example, inputs from other neurons would drive changes in the electrical activity, but not changes in channel expression. On the other hand, if the cell is ‘trying’ to maintain some mean calcium concentration, or some other measure of activity, then changing the environment in which the neuron operates will change channel expression. As an extreme example, if we rip the neuron from its network and put it in a dish, the normal pattern of rhythmic bursting will go (wildly) wrong, but the calcium-sensitive dynamics of the channel expression levels will eventually bring the system back into something close to the original pattern.

⁷⁰ As in the case of kinetic proofreading, I think there is a tendency to remember the original papers as having proposed mechanisms that solve problems. But I think that, in many ways, it was a much deeper contribution to *formulate* the problems. Even if the solutions turn out not to be precisely the ones chosen by Nature, the problems are important.

In this new state, the channels are playing different roles in the dynamics, because the driving forces for ionic current flow are different, but the final pattern of activity is the same. A literal version of this rather dramatic scenario actually works experimentally, as shown in Fig 105.

We have noted already that, in invertebrates such as flies and crabs, neurons have names, numbers and identifiable functions from individual to individual within a species. This discussion of stabilizing patterns of activity rather than expression levels suggests that this reproducibility of function can be achieved without exactly reproducing the number of copies of each channel protein. Further, although the slice through parameter space shown in Fig 104 suggests that the region compatible with normal function is convex, this in fact is not the generic case, and real models often have banana-shaped volumes in parameter space which are consistent with particular patterns of electrical activity. [Look through Goldman et al (2001) & Golowasch et al (2002) to decide on a figure.] Again this is consistent with what one sees experimentally, most impressively in subsequent experiments which measure directly the number of copies of mRNA for several channel types in single cells [recent

refs from Eve's group].

One might worry that we have replaced the tuning of channel copy numbers with a fine tuning of the regulatory mechanisms on all the channels. In fact, it is not plausible that calcium acts directly on expression of genes. More likely is that calcium binds to some protein, and when its binding sites are occupied the protein can act, directly or indirectly, as a transcription factor. Then the fact that all the genes have the same calcium dependence to their steady state values reflects the fact that they are all being regulated by the same calcium binding protein. Exploring this scenario in more detail, one realizes that the kinetics of binding and unbinding of calcium to the sensitive protein can span the time scales of action potentials, bursts, and even the basic rhythm itself. By combining signals from calcium binding proteins with different kinetics [that's a little quick!] one can thus stabilize more subtle details in the pattern of electrical activity. Maybe there is more to say about all this before drawing the lessons. Check most recent papers.

Faced with a model that explains the behavior of cells only when parameters are finely tuned, we become suspicious that we are missing something. One possibility—often the most plausible—is that the model simply is wrong. The models that we have for biological systems are not like the Navier–Stokes equations for fluids or the standard model of particle physics; we have many reasons to doubt that we are simply solving the wrong equations. But the electrical dynamics of neurons are a special case. Our mathematical models of channel dynamics emerged as accurate summaries of a huge body of data, and are nearly exact on the time scales that are experimentally accessible. Rather than rejecting the models, we must conclude that we are missing something, presumably on time scales longer than the experiments that go into characterizing the channel kinetics. In particular, what look like constant parameters must become slow dynamical variables. The simplest implementation of this idea seems to work, and to generate several dramatic experimental predictions which have since been confirmed. Indeed, this theoretical work on the problem of parameter determination has launched a whole subfield of experimental neurobiology, investigating the activity-dependent regulation of the ‘intrinsic’ electrical properties of neurons [be sure there is a ref to recent review].




FIG. 105 Changing intrinsic properties of the STG neurons, from Turrigiano et al (1994). At left, an experiment in which one cell is ripped from the network and placed in isolation. At first (top) the electrical activity shifts from rhythmic bursts to repeated (“tonic”) firing of single action potentials. After two days in culture, the cell is silent but responds to small positive currents with tonic firing; after three days the response consists of bursts not unlike those in the native network environment. At bottom, continuous recordings demonstrate that this switch from tonic firing to bursting can occur within an hour. At right, one hour of stimulation with negative current pulses drives a shift from bursting to tonic firing, which is reversed after one hour of no stimulation. All these changes in activity reflect changes in the numbers of different types of ion channels in the cell membrane, as predicted from the models discussed in the text.

Our understanding of ion channels goes back to the classic papers of Hodgkin and Huxley (1952a–d), still very much worth reading. The series of papers (of which the first really is Hodgkin, Huxley & Katz 1952) describes many ingenious experiments, culminating in a mathematical model which predicts the form and speed of the action potential. Include Hodgkin's summaries—Croonian lecture, plus the one from Pursuit of Nature] For a modern textbook account, see Dayan & Abbott (2001). The Hodgkin–Huxley model is

complicated, so over the years there have been various attempts at simplifying to the point where one can gain analytic insight; for one approach, see Abbott & Kepler (1990) [also FitzHugh & Nagumo].

Abbott & Kepler 1990: Model neurons: From Hodgkin–Huxley to Hopfield. LF Abbott & T Kepler, in *Statistical Mechanics of Neural Networks*, L Garrido, ed, pp 5–18 (Springer–Verlag, Berlin, 1990).

Dayan & Abbott 2001: *Theoretical Neuroscience* P Dayan & LF Abbott (MIT Press, Cambridge, 2001).

Hodgkin et al 1952: Measurement of the current–voltage relations in the membrane of the giant axon of *Loligo*. AL Hodgkin, AF Huxley & B Katz, *J Physiol (Lond)* **117**, 442–448 (1952).

Hodgkin & Huxley 1952a: Currents carried by sodium and potassium ions through the membrane of the giant axon of *Loligo*. AL Hodgkin & AF Huxley, *J Physiol (Lond)* **117**, 449–472 (1952).

Hodgkin & Huxley 1952b: The components of membrane conductance in the giant axon of *Loligo*. AL Hodgkin & AF Huxley, *J Physiol (Lond)* **117**, 473–496 (1952).

Hodgkin & Huxley 1952c: The dual effect of membrane potential on sodium conductance in the giant axon of *Loligo*. AL Hodgkin & AF Huxley, *J Physiol (Lond)* **117**, 497–506 (1952).

Hodgkin & Huxley 1952d: A quantitative description of membrane current and its application to conduction and excitation in nerve. AL Hodgkin & AF Huxley, *J Physiol (Lond)* **117**, 500–544 (1952).

For a modern view of ion channels, see Hille (2001). [add some classic references about resonances etc.] For a detailed discussion of system in which the effective resonance generated by channel kinetics has functional importance, see Wu & Fettilplace (2001). [Need references that survey the richness of ion channel diversity, phosphorylation, splicing variants, etc. Check Laughlin refs re splicing variants in the fly eye.] For one example of this complexity, see Tian et al (2001).

Hille 2001: *Ion Channels of Excitable Membranes, Third Edition*. B Hille (Sinauer, 2001).

Wu & Fettilplace 2001: A developmental model for generating frequency maps in the reptilian and avian cochleas. YC Wu & R Fettilplace, *Biophys J* **70**, 2557–2570 (1996).

Tian et al 2001: Alternative splicing switches potassium channel sensitivity to protein phosphorylation. L Tian, RR Duncan, MS Hammon, LS Coghil, H Wen, R Rusinova, AG Clark, IB Levitan & MJ Shipston, *J Biol Chem* **276**, 7717–7720 (2001).

[Need the list of references for the “how we know HH were right” discussion.]

:

The problem of setting the numbers of each kind of ion channel emerged in attempts to make quantitative models of individual neurons in the stomatogastric ganglion. For a recent overview of the STG, emphasizing its role as a model system for studying network dynamics, see Marder & Bucher (2007). These models reached a very high degree of sophistication, as described in the series of papers by Golowasch & Marder (1992), Buchholtz et al (1992) and Golowasch et al (1992). The basic idea of regulating the number of ion channels via feedback from the electrical activity of the cell was described by LeMasson et al (1993); see Abbott & LeMasson (1993) for a more complete account. Dramatic experimental evidence for “self–tuning” of channel numbers came (quickly) from Turrigiano et al (1994). For feedback mechanisms with sensitive to multiple time scales, see Liu et al (1998).

Buchholtz et al 1992: Mathematical model of an identified stomatogastric ganglion neuron. F Buchholtz, J Golowasch, IR Epstein & E Marder, *J Neurophysiol* **67**, 332–340 (1992).

Golowasch & Marder 1992: Ionic currents of the lateral pyloric neuron of the stomatogastric ganglion of the crab. *J Neurophysiol* **67**, 318–331 (1992).

Golowasch et al 1992: The contribution of individual ionic currents to the activity of a model stomatogastric ganglion neuron. J Golowasch, F Buchholtz, IR Epstein & E Marder, *J Neurophysiol*, **67**, 341–349 (1992).

LeMasson et al 1993: Activity–dependent regulation of conductances in model neurons. G LeMasson, E Marder, & LF Abbott, *Science* **259**, 1915–1917 (1993).

Marder & Bucher 2007: Understanding circuit dynamics using the stomatogastric nervous system of lobsters and crabs. E Marder & D Bucher, *Annu Rev Physiol* **69**, 291–316 (2007).

Abbott & LeMasson 1993: Analysis of neuron models with dynamically regulated conductances. LF Abbott & G LeMasson, *Neural Comp* **5**, 823–842 (1993).

Liu et al 1998: A model neuron with activity–dependent conductances regulated by multiple calcium sensors. Z Liu, J Golowasch, E Marder & LF Abbott, *J Neurosci* **18**, 2309–2320 (1998).

Turrigiano et al 1994: Activity–dependent changes in the intrinsic properties of cultured neurons. G Turrigiano, LF Abbott & E Marder, *Science* **264**, 974–977 (1994).

Need references to second generation of experiments on mRNA levels. Maybe some pointers to work on networks??

C. The states of cells

Cells have internal states. Sometimes these states are expressed in a very obvious way, even to external observers, as when we see the alternating black and white stripes of a zebra. In other cases, the states are hidden, as when a neuron stops responding to a constant external stimulus, but then rebounds when the stimulus is removed; the amplitude of the rebound reflects the initial amplitude of the stimulus, which must have been stored in some internal state, separate from the output. In these two examples, we also see that these internal states can be discrete or continuous. In many cases, the states of cells are known to be encoded by the concentrations of particular, identifiable molecules, and these concentrations in turn reflect a balance of multiple kinetic processes. If we try to transcribe these qualitative ideas simply into quantitative models, we will find that the states of cells depend on parameters. Most obviously, these states will depend on absolute concentrations, and there is a widespread suspicion that absolute concentrations are highly variable, making them poor candidates for the markers of cellular state. More generally, it would seem that, unless we are careful, states will depend sensitively on parameters, providing another example of the problem of fine tuning vs. robustness that we have been discussing.

In this section we will look at the issue of fine tuning in a variety of biochemical and genetic networks. Historically, these discussions have been independent of the earlier work on protein folding or ion channel dynamics, although I hope to make clear that the conceptual questions are the same. We'll start with the problem of adaptation to constant sensory signals, and move to more complex examples in the cell cycle and embryonic development.

When you tie your shoes in the morning, you can feel the pressure against the skin of your foot, but very quickly this sensation dissipates. When you step outside on a bright summer morning, you are aware of the light, but soon everything looks normal, and you would have trouble reporting accurately the absolute light level. These are examples of sensory or perceptual adaptation, in which we gradually become unaware of constant stimuli, while maintaining sensitivity to small changes in these incoming signals. One of the first things discovered when it became possible to record the signals propagating along individual nerve fibers is that this adaptation occurs, at least in part, in the response of the single cells that first convert sensory inputs into electrical signals, as shown in Fig 106. Further, as we have seen in the discussion of bacterial chemotaxis (Section II.B), adaptation occurs even in the sensory systems of single celled organisms. As we will discuss in connection with the problems of information transmission in neural coding (Section IV.C), adaptation can be a rich and complex phenomenon, being driven not just by constant background signals, but also by the statistical structure of fluctuations around this background.

In the simplest case, where adaptation consists of reducing the response to constant signals while maintaining




FIG. 106 The original experiments demonstrating adaptation in the response of single sensory neurons (here from the muscle spindle) exposed to constant stimuli (weights), from Adrian & Zotterman (1926a).




FIG. 107 A schematic of the mechanisms underlying sensory adaptation. The branch which generates fast responses insures that sudden changes in input will be transduced faithfully. The branch with the slower response causes a gradual decay of the output in response to constant inputs. To have truly zero response to constant input requires that the two branches be perfectly balanced.

sensitivity to small transient changes, there is a natural schematic model (Fig 107) in which a rapid positive response to the sensory input is cancelled by a slower negative response. In several systems we can identify the molecular or cellular components that correspond to these different branches, and we will discuss the example of bacterial chemotaxis in detail. For the moment, however, our concern is more general. If adaptation is accomplished through some pathway that is independent of the basic response to incoming stimuli, then the 'gain' of the two pathways are set by independent parameters. If we want the responses to constant inputs to be small, then these two gains must be very similar, so that they nearly cancel. In particular, if we want truly zero response to constants—zero net gain at zero frequency—then the signals passing through the two branches need to cancel exactly, and this seems to require fine tuning of the parameters.

Before saying that we have found a problem, we should examine the precision of cancellation that is actually required. In the example of the fly photoreceptors, discussed in Section I.A, we saw that the system acts as a nearly ideal photon counter up to rates of $\sim 10^5$ photons/s. If the response to a single photon lasts (at its shortest) ~ 10 ms, this means that cell is effectively counting up to ~ 1000 . But, as we noted, single photon responses are on the order of a few millivolts, so if things just add up the voltage across the cell membrane would have to change by several Volts, and this isn't going to happen—something like 90 – 99% of this response needs to be cancelled in order to fit into the available dynamic

range.

In the case of bacterial chemotaxis, we have seen in Section II.B that adaptation is essential for function (see especially Problem **). Because the cell makes decisions based on the time course of concentrations along its trajectory, having a response to constant stimuli would mean that the cell effectively confuses “things are good” for “things are getting better,” and this would impede progress up the gradient of desirable chemicals. Direct measurements of the clockwise vs. counterclockwise rotation of the flagellar motor, as in Fig 47, show that the response to a small step in the concentration of attractant molecules decays to zero, so that adaptation is nearly perfect. Another way of seeing this is if one exposes the cells to concentrations that are exponentially increasing in time, the fraction time the motor spends running clockwise become constant, depending on the rate of exponential increase, rather than rising up to saturation; an example is in Fig 108.

Problem 111: Exponential ramps. Give a problem to work out why Fig 108 makes sense!

If we observe freely swimming bacteria, then we can count the rate at which they initiate tumbles, and see that this also adapts to constant stimuli; Fig 109 shows an unnatural but dramatic example, in which a population of bacteria is suddenly exposed to millimolar concentrations of aspartate, starting from zero background concentration. Tumbling is almost completely suppressed




FIG. 108 Response of *E. coli* to exponentially increasing (top) or decreasing (bottom) concentrations of an attractant, from Block et al (1983). **Probably needs more explanation.**




FIG. 109 Experiments on adaptation in a large population of *E. coli* (Alon et al 1999). At time $t = 0$, the population is exposed to a high concentration of an attractive chemical, and as a result the bacteria almost stop tumbling. Over time, they adapt, and the average rate of tumbling approaches the steady-state value observed in the absence of stimuli.

for nearly ten minutes, but eventually recovers to within $\sim 10\%$ of its initial rate, despite the fact that the initially saturating stimulus continues to be present [show an earlier figure of this flavor from Berg or Koshland?].

To understand how it’s possible to achieve near perfect adaptation without fine tuning of parameters (as one might have thought from Fig 107), we have to dig into the details of the molecular mechanisms involved. In Section II.B we outlined the fast events involved in the “positive” part of the chemotactic response (Fig 48). To review briefly, receptor molecules on the cell surface form a complex with the enzyme CheA (a kinase), held together by a scaffolding molecule CheW. The complex is in equilibrium between the active (CheA*) and inactive (CheA) states, and this equilibrium is shifted by binding of attractant or repellent molecules to their receptors; for attractants, binding shifts the equilibrium toward the inactive state. The active kinase CheA* phosphorylates the protein CheY, which can diffuse through the cell from the receptor complex to the flagellar motor, where it binds and favors clockwise rotation, driving the tumbling motion of the cell; the action of the kinase is opposed by a phosphatase, CheZ. Thus, an increase in the attractant concentration drives the kinase toward its inactive state, reducing the rate of phosphorylation of CheY; the continued action of the phosphatase results in a reduction of the CheY-P concentration, and this reduces the probability of tumbling. This whole pathway is extraordinarily sensitive, responding reliably to individual molecules as they bind to their receptors.

How does the extremely sensitive response of the chemotactic system get cancelled when stimuli are maintained at constant levels? In addition to binding the

chemoattractant or repellent molecules, the receptors can be modified by covalent attachment of methyl groups. Much as with ligand binding, these modifications shift the equilibrium between active and inactive conformations of the kinase CheA—binding of attractants favors the inactive state, addition of methyl groups favors the active state. The key point is that the active kinase not only phosphorylates CheY, leading to clockwise rotation of the motor, it also phosphorylates CheB, and then CheB-P removes methyl groups from the receptor. Thus, when an attractant lowers the activity of the kinase, it also allows more methyl groups to be attached, driving the activity back toward its original level—adapting.

Although the methylation system provides a pathway to cancel the effect of the immediate response to sensory inputs, it isn't clear that this cancellation should be anywhere near exact. In general, one would need to tune the activity of the methylation and de-methylation enzymes to make sure that their effects exactly balance the direct response to sensory input. So, this system provides an example of our general problem of fine tuning, as emphasized by Barkai and Leibler. In addition to identifying the problem, they proposed that one can evade this need for fine tuning by assuming that the de-methylation enzyme CheB only recognizes the active state of the receptor-kinase complex, and ignores the inactive conformation. If this is true, one doesn't even need the phosphorylation of CheB in order to close the feedback loop.

To see how the Barkai-Leibler scheme works, let's




FIG. 110 Methylation of the receptors allows for adaptation of the chemotactic response. At left, addition of methyl groups acts, similarly to ligand binding, as an allosteric effector, shifting the equilibrium between the active and inactive states of the kinase CheA; the schematic is meant to indicate that there are multiple methylation sites. At right, the feedback loop is closed by having the active kinase CheA* trigger activation of the de-methylation enzyme CheB. **Need to redraw to remind that methylation is working opposite to the effects of an attractant binding.**

imagine that the whole receptor complex, which might include a cluster of several receptor molecules, switches as a whole between active and inactive states. There is some free energy difference ΔF between these states, and there are two contributions to this difference—one from the binding of attractants, and one from methylation. Assume that the contribution of the methyl groups is additive, and that the contribution from ligand binding has some arbitrary dependence on ligand concentration c (which we could work out from a model like that in Fig 110; see Problem 112 [check] below). Then the number of active enzymes is given by

$$A^* = \frac{A_{\text{total}}}{1 + \exp[F_L(c) - n_M \Delta_M]}, \quad (552)$$

where n_M is the number of methyl groups per receptor complex. This number reflects a balance between the activities of CheR and CheB, so we can write schematically

$$\frac{dn_M}{dt} = V_R - V_B, \quad (553)$$

where V_R and V_B are the ‘velocities’ of the methylation and de-methylation enzymes, respectively.

The key assumptions suggested by Barkai and Leibler are that CheR is running at some maximal rate, limited by its internal dynamics and not by the availability of substrate, while the velocity of CheB does depend on the availability of its substrate A^* according to some function $f(A^*)$ that we don't need to specify. Then

$$\frac{dn_M}{dt} = V_R^{\text{max}} - V_B^{\text{max}} f(A^*). \quad (554)$$

In order to reach steady state ($dn_M/dt = 0$), we must have

$$A^* = A_0^* = f^{-1}(V_R^{\text{max}}/V_B^{\text{max}}), \quad (555)$$

independent of the ligand concentration c . Thus all steady states in the system must have the same level of activation of the kinase, hence the same level of phosphorylation of CheY and the same rate of tumbling. These steady states at varying c are not identical—they involve different levels of methylation—but they have the same functional output.

Problem 112: Allosteric model for chemotactic receptors. [check for earlier problem about this ...] The schematic in Fig 48 is equivalent to a Monod-Wyman-Changeaux model (all relevant pointers) in which the whole complex of the receptor, CheW and CheA has two states, and the equilibrium is shifted by binding of the attractant molecule. In Fig 110, attachment of methyl groups also shifts this equilibrium, but the binding and unbinding of these groups is part of an energy-yielding reaction, and so doesn't have to obey detailed balance. Show that, nonetheless, these schematics generate Eq (552), which has a decidedly Boltzmann form. Why

does this work? What would change if groups or clusters of N receptor complexes were tied together, and forced to all be in the same activation state?

If the scenario sketched here is correct, then we should be able to test it by manipulating the activity of the methylation and de-methylation enzymes, using the modern tricks of molecular biology to modify the genome of *E. coli*. To begin, one can replace CheB with a mutant form which cannot be phosphorylated; adaptation still works, and still is nearly perfect, suggesting that phosphorylation is not the key step in closing the feedback loop. Then one can delete the normal CheR gene and replace it with a plasmid which carries the CheR coding region under the control of a promoter that responds to external signals. In this way one can generate roughly 100-fold variations in CheR expression levels, from half the normal level to 50× over-expression, as in Fig 111. Throughout this range, adaptation to large inputs (as in Fig 109) is within $\sim 10\%$ of being exact. Although the mean rate of tumbling to which the system adapts, as well as the time scale of this adaptation, depends on the amount of CheR in the cell, the fact that this rate is independent of input concentration does not. **Are there experiments that look at adaptation in response to smaller signals? Maybe from Sourjik?**

There is a lot of evidence that the methylation level of the receptors really is the molecular representation of the cell's adaptation state. As such, we might have expected that over- or under-expressing the enzyme that carries

out the methylation reaction would shift the actual state of the system, and this would show up as a change in the output. In the model considered here, however, this last expectation is violated. The absolute level of kinase activity, and hence the absolute tumbling rate, does indeed change when we change the expression of CheR. But Fig 111 shows us that the average steady state response to an applied step in attractant concentration remains zero, independent of the CheR level. Thus, the precision of the balance between the processes responsible for excitation and adaptation does not depend on fine tuning of the underlying kinetics.

Problem 113: Calcium driven adaptation in neurons.

Consider a neuron that generates spikes at rate r . Let's assume the response to external inputs I involves this rate relaxing toward some steady state,

$$\tau \frac{dr}{dt} = r_{\max} f(I, [\text{Ca}]) - r, \quad (556)$$

where we note explicitly that the rate depends both on the inputs and on the intracellular calcium concentration. Write an equation for the dynamics of $[\text{Ca}]$, assuming that each spike brings in a fixed number of calcium ions, and that there is a pump which extrudes the ions at some opposing rate. The pumping rate must depend on the concentration, but for the moment take this dependence as some unknown function $V_{\text{pump}}([\text{Ca}])$. Find equations that describe the steady state of this system. Are there conditions on $V_{\text{pump}}([\text{Ca}])$ that lead to a steady state spike rate that is independent of the input I ? If the input changes suddenly, does the spike rate still respond? Explain how this relates to the discussion of chemotaxis given here.




FIG. 111 Chemotactic responses in the presence of varying amounts of CheR, from Alon et al (1999). At the top, ‘adaptation precision’ is measured as the ratio of the mean tumbling rates in the presence and absence of 1 mM aspartate (as in the experiments of Fig 109). The actual tumbling rates and the time required to reach steady state after sudden exposure to 1 mM aspartate are shown in the bottom panel.

Are we done? I think there is still more to this problem. To begin, the fact that motor output is an extremely steep function of the CheY-P concentration (**pointer**) means that successful adaptation requires more than just a constant level of CheY-P in steady state, independent of the input signal—this level actually has to fall into a very narrow range, or else the cells will be always running or tumbling. The parameters which determine the steady state level of CheY-P are independent of the properties of the motor, which determine the functional operating range for this concentration. This seems like the same sort of balancing problem that Barkai and Leibler were worried about, but in a different part of the system, where their solution has no obvious analog. **[Has somebody worried about this?]**

Next, you should also be a little suspicious about the simple equations above. At best, they are some sort of mean field theory in a system where fluctuations could be important. Also, while it's plausible that CheB recognizes CheA* as opposed to CheA, one might worry that the rate of removing methyl groups depends on how many are there (especially if that number goes to zero!). There

must be some regime in which the simple argument is right, but we need a more honest calculation. [several groups have tried this; look closely at Wingreen et al, and check for others, to decide what to say here.]

Finally, although one can manipulate the *E coli* genome to change expression level of individual proteins by large factors, the many protein components of the chemotactic system are encoded on just two operons, which means that the expression of the different components is tightly coupled under normal conditions [be sure to have talked about operons before this, or maybe this is really a good place to introduce the idea?].

The *mocha* operon encodes CheA and CheW, along with the flagellar motor proteins, and the *meche* operon encodes CheR, CheB, CheY and CheZ, along with two classes of receptor proteins. Recent experiments indicate that there is covariation even between the expression levels of CheA and CheY, suggesting that the cell can in fact control at least the relative concentrations of these proteins fairly precisely. Further, there is direct evidence that tight correlation between protein concentrations actually improves chemotactic performance, as shown in Fig 112.

Problem 114: Balancing CheY and CheZ. Take the students through a model in which it becomes clear why variations in the relative levels of CheY and CheZ are detrimental for chemotaxis, thus making sense out of Fig 112.

It is interesting to compare the problem of robustness vs. fine tuning in the case of chemotaxis with what we learned in the case of ion channels (Section III.B). For ion channels, function really does depend sensitively on the number of copies of the different proteins in the network, and neurons have evolved control mechanisms that use their functional output (or a near surrogate) to control these copy numbers. Importantly, there are many ways to achieve the same function, so it is not the number of copies of each component that is tuned, but rather some possibly complex combinations of these quantities. For chemotaxis, the message of the experiments in Fig 111 is that large variations in the copy number of just one component can be tolerated, pointing toward networks that are intrinsically insensitive to this parameter variation rather than any hidden control or tuning mechanisms. This suggests that one system is tuned, and the other is robust.

On the other hand, Fig 112 shows that, as with ion channels, the relative copy numbers of the proteins in the chemotaxis network *are* controlled, and that this control contributes to function. Experiments more directly




FIG. 112 Better chemotactic performance is associated with correlated fluctuations in protein levels, from Løvdo et al (2009). At right, *E coli* swarm outward toward attractants. Cells have been engineered to express CheY and CheZ only under the control of a promoter induced by external signals. If we select cells from regions B or C of the swarm, we see that the cells which have been efficient (B) have tightly correlated variations in the two protein levels, while cells that have been inefficient (C) have weaker correlations. Thus, selection for chemotactic efficiency will drive down the *relative* fluctuations in expression levels, even there is substantial tolerance for variation in the absolute levels.

analogous to Fig 111 have now been done in stomatogastric ganglion neurons, and one finds that there are control mechanisms which can compensate for over-expression of particular channel types by changing the expression levels of other channels see Fig 113. Perhaps surprisingly, these compensation mechanisms are triggered even if the first channel is non-functional and hence doesn't effect the electrical output, suggesting that there are signals internal to the transcriptional and translational networks which encode something about the correct, functional operating point of the system. This could be a much more general phenomenon.

Before moving on, there is also a somewhat philosophical point to be made about the mechanism of robustness in chemotaxis, or perhaps even about the idea of robustness itself. If we expect the function of a biochemical network to be robust against parameter variation, this robustness must be a property of the network topology—which nodes (molecules) are connected by arrows (reactions). In the specific model considered by Barkai and Leibler, for example, it is essential that CheB acts on CheA* as a substrate, but not on the inactive CheA—what is important in this case is the *absence* of a link in the network connecting CheB with CheA.

The particular links that appear (or don't appear) in a biochemical network reflect the specificity of the various enzymatic and protein-protein binding reactions. Substrate specificity is a classical topic in biochemistry, and much of what we understand about this topic was learned through painstaking experiments on purified samples of particular enzymes. The ideas of robustness emerged at a time when the community started to wonder if there wasn't something a bit hopeless about the overall project

of this classical biochemical approach. While one can study individual enzymes in detail, many interesting biological functions emerge from networks with many interacting proteins engaged in dozens if not hundreds of individual reactions. Further, the conditions inside the living cell may be far from those that we can reproduce in a test tube. How then could we ever study every one of the relevant reactions, under the right conditions? Seen in this light, it just doesn't seem plausible that the accumulated biochemical knowledge will "add up" to give us an understanding of how cells function as complete systems. Robustness was one of several ideas offered as an alternative—if Nature has selected networks that are robust to parameter variations, then the (already somewhat hopeless) project of measuring all these parameters could safely be abandoned. But because network topology is an encoding of substrate specificity, we can't really brush aside all of classical biochemistry. Indeed, the example brought forward by Barkai and Leibler is one in which the biochemistry is subtle, with one protein recognizing different conformations of another. At the end of the day, then, biochemistry has its revenge—robustness may be an emergent, system level property, driven by network topology, but this topology is an expression of the underlying, detailed biochemistry.

[I'd like put here a discussion of the work by Tang and coworkers on the cell cycle. The idea is that not just states, but trajectories through the space of states, are robust. Need to sort through the papers for details. Should also discuss the results from the Siggia/Cross col-




FIG. 113 Responses to over-expression of a channel, from Maclean et al (2003). At left, injecting mRNA for the A current channel (see Table I) produces, after 72 hrs, an increase in the current that flows when the voltage is stepped through the range expected to activate this channel. This shows that injecting the mRNA really does result in more channels being synthesized and inserted into the membrane. At right, a demonstration that this increased number of channels (in the "Shal" trace) does not perturb the basic pattern of activity (seen in the control). This is possible only because the cell compensates by increasing the expression levels of other channels.

laboration on the cell cycle, in particular what makes the decision to go from one state to the next reliable and irreversible.]

Another example of "robust" output from a biological network is that almost everyone you know was born with five fingers on each hand. In insects, one can count even more instances in which discrete pieces of the body are arranged in a repetitive pattern, from the segments of the body itself (as in the beautiful caterpillar shown in Fig 114) to the hairs or bristles on the body surface; essentially every member of a particular species has the same number of body segments, the same number of hairs, and even the positions of the hairs are identifiable from individual to individual. It is not at all obvious how this level of reproducibility is achieved.



FIG. 114 Insects provide many examples of repeated, reproducible structures visible on the outside of the body. Image of tiger moth caterpillar from <http://www.hsu.edu/content.aspx?id=7435>. [probably should take our own picture; also, maybe another panel about segments, or bristles?]

Broadly speaking we can distinguish two classes of explanations for the reproducibility of pattern formation ('morphogenesis') in the embryo. In the first kind of explanation, the organism works to set the initial conditions and boundary conditions very precisely, and each step in the process has been tuned to minimize noise. Patterns then develop in a reproducible fashion in the same way that accurate clocks continue the same time even though tick independently. Alternatively, it is possible that noise and errors abound, but that there are error correction mechanisms that pull the pattern back to its ideal structure. Of course, it is also possible that both scenarios are correct: nature has selected for systems with minimal noise, and taken care to control the conditions of development, but error correction mechanisms still are needed to deal with the vagaries of a fluctuating environment.

To appreciate why the observations of reproducibility in morphogenesis are so puzzling, we need to review some of the basic mechanisms by which patterns form in the developing embryo. We will also need to check our qualitative impressions of reproducibility against quantitative data. Let's start with the background.

We recall that embryos start as just one cell, the fertilized egg, and then there are multiple cell divisions. Every one of these cells (as in our adult bodies) has the same DNA, assuming that nothing has gone wrong. What makes the different cells different is that they “express” different genes. The genes code for proteins, but not all of the proteins are made in all cells; the reading of the code to make the proteins is called the expression of the genes, as we have discussed before. Importantly, the regulation of gene expression is not just the flipping of a switch sometime in development, but rather something that all cells (from neurons in our brain down to bacteria) are doing all the time. Embryos come in all shapes and sizes throughout the animal kingdom, but for various reasons people have focused on a few model systems, and we will do the same.

The fly embryo is an interesting model system for many reasons. One is that there is a well developed genetics for fruit flies (the species *Drosophila melanogaster*), made possible not least by their rapid growth and reproduction. Embryonic development itself is rapid as well, leading from a fertilized egg to the hatching of a fully functional maggot (the larvae of flies, like caterpillars for butterflies) within 24 hours. All of this happens inside an egg shell, so there is no growth—pattern formation occurs at constant volume. The egg is $\sim 1/2$ mm long, so one starts with one rather large cell, which has one nucleus. In the maggot there are $\sim 50,000$ cells. For the first three hours of development, during which the “blueprint” for the body plan is laid out, something special happens: the nuclei multiply without building walls to make separated cells. Thus, for about three hours,

the fly embryo is close to the physicists idealization of a box in which chemical reactions are occurring, with the different molecules free to move from one end of the box to the other (perhaps even by diffusion, although this is a more subtle question).

The duplication of the nuclei is more or less synchronous for the first 13 mitotic divisions, or nuclear cycles, which is visually quite striking. During cycles 8 through 10, almost all of the nuclei move to the surface of the egg, where they form a fairly regular two dimensional lattice; conveniently, with all the nuclei at the surface of the egg, we have a much better chance to “see” what is going on (see Figs ?? et seq). With each subsequent cycle, this lattice dissolves and reforms. With cycle 14, the synchronous duplication of nuclei stops, and there is a pause while the embryo builds walls between the nuclei to make separate cells. If you stop the action at this point and take an electron micrograph of the embryo, what you see is at the top in Fig 115. If you count, you’ll find that there are ~ 6000 cells on the surface. This is smaller than 2^{13} , but that’s because not all of the nuclei make it to the surface; some stay in the interior of the embryo, probably not by accident since these become cells with special functions. Notice that all the cells look pretty much alike. If instead of stopping at this point, we wait just 15 minutes more, we see something very different, shown at the bottom in Fig 115. Notice that there is a vertical cleft, about one-third of the way from the left edge of the embryo. This is the “cephalic furrow,” and defines which part of the body will become the head. There is also a furrow along the bottom of the embryo, which is where the one layer of cells on the surface starts to fold in on itself so that you can have two “outside” surfaces (think about the inside and outside of your cheek, both of which are outside of the body from the topological point of view—we are not simply connected!), a process called “gastrulation.”

It’s not just that the embryo breaks into a head and a non-head. In fact there are many different pieces to the body, usually called segments, as noted above. The obvious question is how the cells at different points in the embryo know to become parts of different segments. The answer is quite striking, and one of the great triumphs of modern biology. Long before cells start moving around and making the three dimensional structures that one sees in the fully developed organism, there is a “blueprint” that can be made visible by asking about the expression levels of particular genes. A now classic set of genetic experiments showed that the number of genes that are relevant in these early patterning events is small, on the order of 100 out of the roughly 25,000 genes in the whole fly genome; if we focus on the pattern along the anterior–posterior axis of the embryo, the number of relevant genes is less than 20. Most of these genes code for transcription factors that control the expression of other genes.




FIG. 115 Electron micrographs of a *Drosophila* embryo in cycle 14, before (top) and after (bottom) gastrulation. Note, in particular, the cephalic furrow roughly one third of the distance from the left in the bottom image. Micrographs taken by EF Wieschaus.




FIG. 116 [...] Thanks again to EF Wieschaus for these images.

Suppose we stop that action in the embryo at cycle 14 and measure the concentration of two of these key proteins. One way to do this is to make antibodies against the protein we are interested in, and then make antibodies against the antibodies, but before using the secondary antibodies we attach to them a fluorescent dye molecule. So if we expose the embryo first to one antibody (which should stick to the protein we are interested in, and not anywhere else, if we're lucky) and to the other, we should have the effect of attaching fluorescent dyes to the protein we are looking for, and hence if we look under a microscope the brightness of the fluorescence should indicate the concentration of the protein (not obvious if this relationship is quantitative; hold that question). One such experiment is shown in Fig 116. Evidently the concentration of the proteins varies with position, and this variation corresponds to a striped pattern. The stripes should




FIG. 117 A combination of Figs 115 and 116, emphasizing that the cephalic furrow occurs along a single line of cells that can be identified from the pattern of pair rule gene expression.





FIG. 118 Antibody staining for the protein Bicoid in the early *Drosophila* embryo, from the original experiments by Driever & Nüsslein-Vollhard (1988a). The plot at the bottom represents means and $2\times$ standard deviations from ten embryos; units of staining intensity are arbitrary. [Are these error bars really $2\times$ the standard deviation, or just \pm the standard deviation? Might have to ask the authors—for a preliminary result, one can't complain about a factor of two, but I want to get it right!]

remind you of the segments in the fully developed animal, and this is actually quite precise. Mutations that move the stripes around, or delete particular stripes, have the expected correlates in the pattern of segmentation. To illustrate this point, we can blow up corresponding pieces of this image and the electron micrograph above, showing the cephalic furrow (Fig 117); hopefully you can see how the furrow occurs at a place defined by the locations of the green and orange stripes. At the moment the names of these molecules don't really matter. What is important is to realize that the macroscopic structure of the fully developed organism largely follows a blueprint laid out within about three hours after fertilization, and that this blueprint is "written" as variations in the expression level of different genes. Furthermore, we know which genes are the important ones, and there aren't too many of them.

We have pushed the problem of pattern formation in the embryo back to spatial variations in the pattern of gene expression, but how do these arise? You could imagine, as Turing did, that these patterns reflect a spontaneous breaking of symmetries in the egg. This, for better or worse, is not how it works. When the mother makes the egg, she places the mRNA for a handful of proteins at cardinal points. For example, there is a protein called Bicoid for which the mRNA is placed at the end that will become the head; importantly, the mRNA is attached to the end of the egg, not free to move. Once the egg is laid, translation of this mRNA begins, and the resulting Bicoid protein is free to move through the embryo. If we

use the same trick as above and stop the action, labeling the embryo antibodies against the protein, we see images like those in Fig 118. Evidently there is a rather smooth gradient in the concentration of Bicoid protein, high at one end and low at the other. A cell sitting at some point in the embryo thus can “know” where it is along this long (anterior–posterior) axis by measuring the Bicoid concentration. This is an example of the very general idea of “positional information” in the embryo.

Since Bicoid is a transcription factor, it provides an input signal to a whole network of interacting genes, and this network can (if we speak colloquially) interpret the positional information, ultimately driving the emergence of the beautiful striped patterns as in Fig 117. We’ll look in more detail at how this happens, but for now let’s try to sharpen our questions about reproducibility.

Measurements on the profile of Bcd concentration show rather decent agreement with an exponential decay, as was noted already in the very first experiments (Fig 118), so that

$$c(x) \approx c_0 e^{-x/\lambda}, \quad (557)$$

where x is measured from the anterior end of the egg. Suppose, then, that the cephalic furrow is placed at the point where the Bcd concentration reaches some thresh-

old value θ_{cf} . The position of the cephalic furrow is then

$$x_{cf} = \lambda \ln(c_0/\theta_{cf}). \quad (558)$$

Thus, if c_0 changes by $\sim 10\%$, the location of the furrow would shift by $\delta x_{cf} \sim 0.1\lambda$. Experimentally, modern experiments show that $\lambda \sim 100 \mu\text{m}$, and the location of the cephalic furrow is reproducible with a standard deviation of $\sim 1\%$ of the length of the embryo, or $\sim 5 \mu\text{m}$ in absolute length. In fact, one can look at other positional markers, such as the locations of peaks or troughs in the striped patterns of expression for the “pair rule” genes in Fig 116, and these are all reproducible at the $\sim 1\%$ level, as shown in Fig 119. Thus, taken at face value, if the Bcd profile provides the basic “map” of position along the anterior–posterior axis, then the absolute concentration of Bcd, c_0 , would have to be reproducible to better than $\sim 10\%$ from embryo to embryo in order to generate the observed reproducibility of these patterns. This problem exists even before we ask how to maintain constant proportions in the face of variations in the overall size of the embryo.

I think that, when people started to think about this problem quantitatively, it seemed implausible that the reproducibility of embryonic development would depend on controlling absolute concentrations with 10% accuracy. On the other hand, the paper which first characterized the spatial profile of Bicoid protein actually reported data on the variations across embryos (Fig 118), and the results are roughly consistent with $\sim 10\%$ reproducibility, at least near the anterior end of the embryo.

Let’s take seriously the simplest possible model for the spatial profile of Bicoid (Bcd), described above in words: the mRNA placed by the mother acts (as it is translated) as a source, the Bcd protein diffuses through the embryo, and the protein is also degraded by some first order reaction. If we simplify and think of the system as just being one dimensional (along the anterior–posterior axis), then the concentration $c(x, t)$ should obey

$$\frac{\partial c(x, t)}{\partial t} = D \frac{\partial^2 c(x, t)}{\partial x^2} - \frac{1}{\tau} c(x, t), \quad (559)$$

where τ is the lifetime of the protein against degradation. The boundary conditions are

$$-D \frac{\partial c(x, t)}{\partial x} \Big|_{x=0} = R, \quad (560)$$

$$\frac{\partial c(x, t)}{\partial x} \Big|_{x=L} = 0, \quad (561)$$

where R is the strength of the source at $x = 0$ and the last condition states that there is no flux out of the other end of the embryo. If we imagine that development is slow enough for the system to come to steady state, and that the embryo is long, the concentration profile becomes

$$c_s(x) = \frac{R\tau}{\lambda} e^{-x/\lambda}, \quad \lambda = \sqrt{D\tau}. \quad (562)$$




FIG. 119 Reproducibility of various spatial markers along the anterior–posterior axis in the early *Drosophila* embryo, from Dubuis et al (2011). At top, fluorescent antibody staining of the protein Eve; scale bar is $50 \mu\text{m}$. Middle, normalized spatial profiles of the fluorescent intensity in 14 embryos; the darker red line is the embryo shown at the top. At right, a small region is blown up to show the variability of the peak; error bars show standard deviations of position and amplitude. Bottom, standard deviations of position for peaks and troughs of several gene expression profiles, as well as for the position of the cephalic furrow measured in live embryos.

Problem 115: Details of the Bicoid profile.

(a.) What are the units of concentration in one dimension? Show that, with this proper choice of units, R is the number of Bcd molecules being translated per second.

(b.) Derive the steady state solution in Eq (562). What is the precise criterion for the embryo to be long enough that this is approximation is accurate?

(c.) At this writing, there is controversy about whether the Bcd profile really reaches steady state during the early stages of development. Although this is an experimental question, we can ask what the simplest model predicts. Intuitively, there is some time scale t_* do you expect the solution of Eq(559) reaches steady state; how does this time scale relate to the other parameters of the problem? Answer this without doing any detailed calculations, and think about how your intuition might go astray.

(d.) Try to do a more detailed calculation to address the approach to steady state. It is useful to assume from the beginning that L is large [in the sense of part (b.)], and to replace the boundary condition at $x = 0$ with a source in the symmetrized version of the problem,

$$\frac{\partial c(x, t)}{\partial t} = D \frac{\partial^2 c(x, t)}{\partial x^2} - \frac{1}{\tau} c(x, t) + 2R\delta(x), \quad (563)$$

where now $-\infty < x < \infty$; be sure you understand why we need a factor of 2 in front of the source term. At $t = 0$, before any protein has been translated, we must have $c = 0$ everywhere. By Fourier transforming in space, show that the exact time dependent solution is

$$c(x, t) = 2R \int_{-\infty}^{\infty} \frac{dk}{2\pi} \frac{e^{ikx}}{Dk^2 + 1/\tau} \left[1 - e^{-(Dk^2 + 1/\tau)t} \right]. \quad (564)$$

Verify that this approaches $c_s(x)$ from Eq (562) as $t \rightarrow \infty$.

(e.) Find a simple closed form for the time derivative of concentration at a point, $\partial_t c(x, t)$. Show that, expressed as a fraction of the local steady state concentration, this derivative peaks at a point $x_* = 2\lambda t/\tau$, and that at this peak $[\partial_t c(x_*, t)]/c_s(x_*) = 1/\sqrt{\pi\tau t}$.

(f.) Suppose we could establish experimentally that, for example, after $t = 1$ hr, at each point x that we can see, $c(x, t)$ changes by less than 1%/min (or $\sim 10\%$ across the time required for the cell cycle). What can you conclude about the parameters of the system, taking the simple model seriously?

We see that this simplest model recovers Eq (557), which was suggested by the data. It gives us an explicit formula for the length constant λ , and tells us (not surprisingly) that the absolute concentration scale c_0 is proportional to the strength of the source—that is, to the rate at which proteins can be translated from the mRNA bound to the anterior end of the embryo. In this simple model, then, if we want c_0 to be reproducible with 10% accuracy, the source strength must also be reproducible. Is it plausible that the mother can count out mRNA molecules, with 10% accuracy, and create an environment in the embryo where the efficiency of translation is similarly well controlled? Alternatively, can we escape from these requirement of fine tuning by moving away from the simplest model?

Suppose that the processes which degrade the Bicoid molecule act not on individual molecules, but on dimers,

and these dimers are rare. We then expect that the concentration of dimers will be proportional to the square of the Bcd concentration, and the dynamics become [instead of Eq (559)]

$$\frac{\partial c(x, t)}{\partial t} = D \frac{\partial^2 c(x, t)}{\partial x^2} - \frac{1}{\tau c_2} c^2(x, t), \quad (565)$$

where c_2 is the concentration scale for dimer formation. Now the steady state solution must obey

$$\frac{d^2 c_s(x)}{dx^2} = \frac{1}{D\tau c_2} c_s^2(x). \quad (566)$$

Notice that if we look for a solution of the form $c_s(x) = Ax^n$, we have

$$\frac{d^2(Ax^n)}{dx^2} = \frac{1}{D\tau c_2} (Ax^n)^2 \quad (567)$$

$$An(n-1)x^{n-2} = \frac{A^2}{D\tau c_2} x^{2n}, \quad (568)$$

which is solved by $n = -2$ and $A = 6D\tau c_2$. Thus, far from the source, the concentration profile is $c_s(x) = 6D\tau c_2/x^2$ independent of the strength of the source. More precisely, to match the boundary condition describing the source at $x = 0$, we have to have

$$c_s(x) = \frac{6D\tau c_2}{(x + x_0)^2}, \quad x_0 = (12D^2\tau c_2/R)^{1/3}. \quad (569)$$

The strength of the source appears only in x_0 ; for $x \gg x_0$ this term is negligible, and for large R this condition itself sets in at very small x . In this model, then, just making the source very strong—but not setting the strength precisely—is sufficient to insure that almost the entire concentration profile will be independent of variations in this source strength.

Problem 116: Fill in the arguments leading to Eq (569).

It is interesting that a relatively small change in molecular mechanism makes such a dramatic change in the robustness of the system to variations in parameters. One might object, of course, that here is no free lunch here. While Eq (569) predicts that the Bcd profile is independent of the source strength, the concentration scale is now set by c_2 , which has something to do with the dimerization of the molecules. The source strength R depends on how many copies of mRNA the mother places in the egg, but the scale c_2 is determined by more global physical-chemical parameters of the cytoplasm, and perhaps these are easier to control. On the other hand, if degradation is

active, via enzymatic reactions, then τ itself will depend on the number of copies of the enzyme that are present in the embryo. Still, it is interesting to ask whether Nature makes use of such a scheme to reduce the sensitivity of morphogen profiles to variations in the strength of the source. [Should also discuss Bollenbach et al (2005).]

Another approach is to give up on making a single morphogen signal reproducible, and to assume that the embryo makes use of multiple signals, hoping that the dominant sources of variation are in a “common mode” that can be rejected by the network that processes these signals. Several models of this flavor have been suggested [refs: Houchmandzadeh et al (2005), McHale et al (2006), others?]. Need to sort out how much of this is about scaling, and how much about reproducibility.

Problem 117: [Should be able to get one or two problems from the model in the last two paragraphs!]

With all this theoretical background, what can we say about the experimental situation? As noted at the outset, there are hints from the earliest literature that Bicoid profiles in *Drosophila* might indeed be reproducible. We also know that the notion of robustness should not be exaggerated. The success of classical genetics in identifying the components of these networks immediately tells us that the system is *not* resistant to the elimination of single components. More subtly, one of the key experiments in establishing that Bicoid is a primary source of positional information was to change the number of copies of the *bcd* gene; with more (or fewer) copies of the gene in its genome, the mother makes more (or less) mRNA and hence drives the strength of the Bcd source up (or down). In response to these changes, the patterns in the early embryo shift, with the cephalic furrow in particular moving—with higher concentrations of Bicoid, the embryo tries to make a larger head, as shown in Fig 120.⁷¹ These results suggest that the embryo does not engage mechanisms which buffer the Bcd profile against variations in the strength of the source. For the morphogens whose concentration profile varies along the other axis of the embryo, however, there are signatures of the nonlinear degradation mechanism which, as we have seen, can

⁷¹ It is not so easy to interpret these results quantitatively, because we don't really know if adding more copies of the gene produces *proportionately* higher concentrations of the protein. Still, Fig 120 is *prima facie* evidence against robustness of the pattern to variations in the strength of the source.




FIG. 120 Variations in the position of spatial markers along the anterior–posterior axis in response to changes in the number of copies of the *bcd* gene, from Driever & Nüsslein–Vollhard (1988b). [explain!]

generate substantial robustness. [Should we have a figure from Eldar et al?]

If there is no buffering, then it really does seem that reproducible outputs require reproducible inputs. Can we see this directly? As discussed in Section II.B, one can genetically engineer flies to express a fusion of Bcd with the green fluorescent protein (GFP), and show that this fusion protein quantitatively replaces the function of the native molecule. Figure 121 shows measurements of the concentration of Bicoid in nuclei from 15 different embryos, using this Bcd–GFP fusion. The raw fluorescence intensity (or the inferred concentration) is plotted vs. position along the anterior–posterior axis for each nucleus. Evidently the variability from embryo to embryo is small, with a standard deviation of less than 20%, and some of this variability can be traced to measurement errors, suggesting that the true variability is $\sim 10\%$ or even less. If the mother has only one copy of the Bcd–GFP gene instead of the usual two, the fluorescence really is cut in half, so again there is no evidence of mechanisms which buffer the observable profile against variations in the strength of the source. This strongly suggests that the mother can place a reproducible number of mRNA molecules into the egg, and that the apparatus for translation has an efficiency that is constant from embryo to embryo as well. It would be attractive to have direct measurements that confirm these conclusions. Of course, this also pushes the problem back. How does the mother count mRNA molecules with $\sim 10\%$ accuracy? How does the embryo ensure that the efficiency of translation, which depends on myriad factors, is reproducible?

Can we make the same argument in any other system? Maybe the Dpp experiments of Bollenbach et al (2008)? Others?

The problem we have been discussing thus far emerges as soon as we claim that position in the embryo is encoded by the concentration of specific molecules. In such a scheme, if we want neighboring cells to do different things, reliably, then we will be driven to questions about how these cells can distinguish small differences in concentration, as discussed in Section II.B.⁷² Conversely, if we want two cells that occupy corresponding positions in different embryos to do the same thing, then we are driven to ask how the concentrations at these corresponding points can be the same. These issues of precision and reproducibility arise even if the size of the embryo and the external conditions of development are identical. There is another problem, related to the variations in size of the embryo, and this is the problem of scaling.

To a remarkable extent, the *proportions* of organisms are constant, despite size variations. We all know people who have especially large heads, but certainly the proportions of the body vary much less than the overall size, and again insects provide clear examples of this, both within species and across species. Different species of flies, for example, have embryos that span a factor of five or more in length, yet they have the same number of




FIG. 121 Measurements of the Bicoid concentration in nuclei along the anterior–posterior axis of the *Drosophila* embryo, from Gregor et al (2007b). Each point corresponds to one nucleus in one embryo; points of the same color come from the same embryo, and error bars show the means and standard deviations across the fifteen embryos in the experiment. The vertical axis shows the fluorescence signal in embryos engineered to make the Bcd–GFP fusion protein, which can be calibrated to give the absolute concentration (at left). The horizontal axis shows the position of the nucleus as a fraction of the overall length of the embryo.

⁷² See also the discussion of positional information, in bits, in Section IV.A.




FIG. 122 Immunofluorescence stainings for products of the gap and pair–rule genes in flies of different sizes, from Gregor et al (2005). (A) Staining of *L sericata* (upper embryos) and *D melanogaster* (lower embryos) for Hunchback (green) and Giant (red) in the left column, and for Paired (green) and Runt (red) in the right column. (B) Staining of *D melanogaster* (upper embryos) and *D busckii* (lower embryo) for Hunchback (green) and Runt (red). Scale bars: 100 μm . [Should give typical sizes of the embryos in the different species!]

body segments, and individual segments have dimensions that scale with the overall size of the organism. You can see this scaling not just in the macroscopic patterns of the developed organisms, but also in the patterns of gene expression, as shown in Fig 122. Indeed, when we have looked at the problem of reproducibility above, we have implicitly used the idea of scaling, always plotting position as a fractional distance along the anterior–posterior axis.⁷³

Scaling is deeply puzzling, perhaps more so for physicists who have thought about pattern formation in non–biological systems. To make this point, let’s imagine making a model of the whole network of interactions that lead to, for example, the beautiful stripes of gene expression. In each nucleus there are chemical reactions corresponding to the transcription of the relevant genes, and the rates of these reactions are determined by the concentrations of the appropriate transcription factors. More equations will be needed to describe translation (although maybe one can simplify, if, for example, mRNA molecules degrade quickly and proteins live longer). Different points in space are coupled, presumably through diffusion of all these molecules, although we should worry

⁷³ [I want to emphasize the distinction between the problems of reproducibility and scaling, but need to think about how to do this. For example, in Fig 121 the embryos have lengths with standard deviation of only $\sim 4\%$.]

about whether diffusion is the correct description. Even if you're not sure about the details, you can see the *form* of the equations: some sort of partial differential equations, in which the local time dependence of concentrations has contributions both from nonlinear terms describing the various chemical reactions and from spatial derivatives describing diffusion or other transport processes,

$$\frac{\partial g_i(x, t)}{\partial t} = D_i g_i(x, t) + F_i(\{g_j\}). \quad (570)$$

But we have seen equations like these before in the study of non-biological pattern forming systems such as Rayleigh–Bernard convection, directional solidification, ...

Many non-biological pattern forming systems generate periodic spatial patterns that remind us of the segments in the insect and the patterns of pair rule gene expression. The scale of these patterns, however, is set by combinations of parameters in the equations. For example, we can combine a diffusion constant with a reaction rate or lifetime to get a length, as in the discussion of the Bicoid profile above ($\lambda = \sqrt{D\tau}$). What happens if you put these equations in a larger box? Well, from Rayleigh–Bernard convection, we know the answer. [should really have an image of convection, or some other ‘physical’ pattern formation problem]. In this system—a fluid layer heated from below—we see a collection of convective rolls, sometimes in stripes and sometimes in 2D cellular patterns. Again, the length scale of the stripes is determined by the parameters of the equation(s). If you put the whole system in a bigger box, you get more stripes, not wider stripes.

Problem 118: A lightning review of pattern formation. [give the students a tour of instabilities etc in some simple case!]

The results in Fig 122 come close to saying that we can put all the same equations into a bigger box, and the stripes come out wider in proportion to the length of the box. One might worry that these are different organisms, and so perhaps evolution has tuned the properties of the proteins involved so that the relevant combinations of parameters turn out to scale with embryo size. The differences can't be too large, because we can identify the same molecules as being involved through similarities of amino acid sequence, and because the same antibodies react with these molecules in different species. Still, it is possible that scaling across embryos in different species reflects an evolutionary adaptation.

If we look across related species of flies with embryos of very different sizes, then the Bcd profiles (as measured

with antibody staining) seem to scale with the length of the egg. One can use the same experimental methods used in making the Bcd–GFP fusion more aggressively, extracting the sequences of Bicoid from flies of different sizes and re-inserting green versions of these different Bicoids into the *Drosophila* genome. The striking result is that the resulting spatial profiles are those appropriate to the host embryo, not the source of the Bicoid. Taken together, all of these results suggest that, as with the problem of variability, the scaling problem is solved at the level of Bicoid itself. It would appear that there is something about the environment or geometry of the embryo itself that couples the global changes in the size of the embryo to the local dynamics.

Scaling might not be so mysterious. Suppose that we think of the (roughly ellipsoidal) embryo as a cylinder, with the source covering one end of this cylinder; since most of the interior of the egg is yolk, we imagine that all degradation of proteins occurs near the surface. If the degradation reaction is rapid, then the surface of the embryo acts as a sink, and in the interior of the embryo the concentration obeys the diffusion equation, with no additional terms. Assuming cylindrical symmetry, the steady state profile must then obey

$$\frac{\partial^2 c_s(x, r)}{\partial x^2} + \frac{1}{r} \frac{\partial}{\partial r} \left[r \frac{\partial c_s(x, r)}{\partial r} \right] = 0 \quad (571)$$

$$-D \frac{\partial c_s(x, r)}{\partial x} = \frac{R}{\pi r_0^2} \quad (572)$$

$$c(x, r = r_0) = 0, \quad (573)$$

where x measures position along the axis of the cylinder (the anterior–posterior axis of the embryo), r_0 is the radius of the cylinder, R is once again the number of molecules per second being injected by the source, and the last condition follows in the limit that degradation reactions at the surface are fast. If we use the standard separation of variables method and look for solutions of the form $c_s(x, r) = e^{-x/\lambda} f(r)$, then we must have

$$\frac{1}{r} \frac{d}{dr} \left[r \frac{df(r)}{dr} \right] + \frac{1}{\lambda^2} f(r) = 0, \quad (574)$$

with the boundary condition $f(r_0) = 0$. You may recognize this as the differential equation which defines the Bessel function,

$$\frac{d^2 J_0(r)}{dr^2} + \frac{1}{r} \frac{dJ_0(r)}{dr} + J_0(r) = 0, \quad (575)$$

so that $f(r) \propto J_0(r/\lambda)$. But then to obey the boundary condition at the surface of the cylinder, we must have $\lambda = r_0/z_{01}$, where z_{01} is the location of the first point where $J_0(z) = 0$. So, in this model, the length scale of the Bicoid profile λ is automatically proportional to the radius of the embryo; if variations in aspect ratio are smaller than variations in length, this will serve, at least approximately, to scale the profile to the size of the egg.

Problem 119: Could it be so simple?. Work out the details of the scenario in the last paragraph. Looking at images of the fly embryo earlier in this section, estimate the radius r_0 assuming that the length of the embryo is $L \sim 0.5$ mm. Does the prediction $\lambda = r_0/z_{01}$ actually work quantitatively?

While simple geometrical mechanisms of scaling might be too simple to work, we should note that embryos of different linear dimensions have the same number of cells. Further, because the nuclei arrange themselves more or less regularly over the embryo surface, the distance from one nucleus to the next provides a local measure proportional to the global size of the egg. **Finish this discussion!!**

Complementary to the problem of scaling is the problem of size control. In many developmental problems (even in later fly development), tissues are growing as they differentiate, and cells have to know both where they are and whether they should still be dividing and hence expanding the size of the tissue. **[Add discussion of work by Shraiman on size control, and subsequent experiments.]**

The discussion so far has taken very seriously the idea that there are “primary morphogens,” placed by the mother, which define provide the basic signal for position in the embryo. Position is a continuous variable, as is concentration. A very different perspective emphasizes that, when development is finished, cells have adopted distinct types or “fates,” which define their function in the adult organism. These fates persist long after the primary morphogen signals have disappeared, and so they must represent stable states of the cells, thus bringing us back to the theme of this section. Cells even maintain their identity and state when separated from their neighbors, which suggests that the biochemical and genetic networks in each cell have multiple attractors. A minimal model of the networks relevant for development, then, would have the right number of attractors but a limited number of dynamical variables, perhaps much fewer than the number of genes involved in the entire network. As with the attractors in the Hopfield model, there is a plausible path to “robustness,” because changing the qualitative behavior of the system would actually require changing the number of attractors—the development of cells into types becomes a matter of topology rather than geometry in the model, and hence invariant to a finite range of parameter variation.

Need to fill out the discussion of attractors. In some ways this is a mathematization of Waddington’s “canalization,” which is an old idea. In modern times, there is work by Reinitz, Sharp and colleagues that tries to make a more direct analogy between genetic and neural

networks. Most recently there is work by Siggia and Carlson on vulva development in *C elegans* that pushes the “minimal model” strategy the furthest, arguing that we can choose coordinates to make the attractors obvious, and then try to map the known biochemical signals into these coordinates, rather than the more usual effort to use biochemical coordinates and decipher the attractors. This belongs here, but isn’t published yet .. hopefully by the time I finalize the text there will be something to cite.

This section needs a conclusion. We have covered a lot of territory, from chemotaxis to development ... what have we learned?

Some of the basic idea about adaptation in sensory neurons were established early on, by Adrian and Zotterman; for a review see Rieke et al (1997).

Adrian 1926: The impulses produced by sensory nerve endings: Part I. ED Adrian, *J Physiol (Lond)* **61**, 49–72 (1926).

Adrian & Zotterman 1926a: ED Adrian & Y Zotterman, The impulses produced by sensory nerve endings: Part II. The response of a single end organ. *J Physiol (Lond)* **61**, 151–171 (1926).

Adrian & Zotterman 1926b: ED Adrian & Y Zotterman, The impulses produced by sensory nerve endings: Part III. Impulses set up by touch and pressure. *J Physiol (Lond)* **61**, 465–483 (1926).

Rieke et al 1997: *Spikes: Exploring the Neural Code*. F Rieke, D Warland, R de Ruyter van Steveninck & W Bialek (MIT Press, Cambridge, 1997).

[Need to check on references for adaptation in bacterial chemotaxis in Chapter 2] Renewed interest in this system was triggered by the work of Barkai and Leibler (1997), who used adaptation in chemotaxis as an example for the more general problem of robustness. The idea that perfect adaptation could be achieved even in the presence of variations in protein copy numbers was then tested more directly by Alon et al (1999). **[Need to reference subsequent work that goes beyond the mean–field level, e.g. from Wingreen et al]** Recent work suggests that, although the biochemical network responsible for chemotaxis may allow for robustness against variations in protein copy numbers, under natural conditions there is relatively precise control over (at least) relative copy numbers, even for proteins on different operons (Kollman et al 2005). In competition experiments, one can even show that tight correlations between protein concentrations improves chemotactic performance (Løvdok et al 2009).

Alon et al 1999: Robustness in bacterial chemotaxis. U Alon, MG Surette, N Barkai & S Leibler, *Nature* **397**, 168–171 (1999).

Barkai & Leibler 1997: Robustness in simple biochemical networks. N Barkai & S Leibler, *Nature* **387**, 913–917 (1997).

Kollman et al 2005: Design principles of a bacterial signalling network. M Kollmann, L Løvdok, K Bartholome, J Timmer & V Sourjik *Nature* **438**, 504–507 (2005).

Løvdok et al 2009: Role of translational coupling in robustness of bacterial chemotaxis pathway. L Løvdok, K Bentele, N Vladimirov, A Müller, FS Pop, D Lebedz, M Kollmann & V Sourjik, *PLoS Biology* **7**, e1000171 (2009).

Pointers toward work on the cell cycle.

Bean et al 2006: Coherence and timing of cell cycle start examined at single-cell resolution. JM Bean, ED Siggia & FR Cross, *Mol Cell* **21**, 3–14 (2006).

Di Talia et al 2007: The effects of molecular noise and size control on variability in the budding yeast cell cycle. S Di Talia, JM Skotheim, JM Bean, ED Siggia & FR Cross, *Nature* **448**, 947–952 (2007).

Li et al 2004: The yeast cell cycle network is robustly designed. F Li, Y Lu, T Long, Q Ouyang & C Tang, *Proc Nat'l Acad Sci (USA)* **101**, 4781 (2004).

Zhang et al 2006: A stochastic model of the yeast cell cycle network. Y Zhang, M Qian, Q Ouyang, M Deng, F Li & C Tang, *Physica D* **219**, 35 (2006).

Lau et al 2007: Function constrains network architecture and dynamics: A case study on the yeast cell cycle Boolean network. K Lau, S Ganguli & C Tang, *Phys Rev E* **75**, 051907 (2007).

Skotheim et al 2008: Positive feedback of G1 cyclins ensures coherent cell cycle entry. JM Skotheim, S Di Talia, ED Siggia & FR Cross, *Nature* **454**, 291–297 (2008).

A modern textbook account of development in the fly embryo is provided by Lawrence (1992). We know which genes are relevant to the earliest events in patterning because of pioneering experiments first by EB Lewis and then by EF Wieschaus and C Nüsslein-Vollhard. Lewis identified a series of puzzling mutant flies where a mutation in a single gene could generate flies that were missing segments, or had extra segments. It is as if the “program” of embryonic development has subroutines (!). Wieschaus and Nüsslein-Vollhard decided to search for all the genes such that mutations in those genes would perturb the formation of spatial structure in the embryo, and they found that there are surprisingly few such genes, on the order of 100. To get a feeling for all this, one can certainly do worse than to read the Nobel lectures from 1994 (Lewis 1997; Nüsslein-Vollhard 1997; Wieschaus 1997).

Lawrence 1992: *The Making of a Fly: The Genetics of Animal Design* PA Lawrence (Blackwell, Oxford, 1992).

Lewis 1995: The bithorax complex: The first fifty years. EB Lewis, in *Nobel Lectures, Medicine or Physiology 1991–1995* N Ringertz, ed, pp 247–272 (World Scientific, Singapore, 1997).

Nüsslein-Vollhard 1997: The identification of genes controlling development in flies and fishes. C Nüsslein-Vollhard, in *Nobel Lectures, Medicine or Physiology 1991–1995* N Ringertz, ed, pp 285–306 (World Scientific, Singapore, 1997).

Wieschaus 1997: Molecular patterns to morphogenesis: The lessons from *Drosophila*. EF Wieschaus, in *Nobel Lectures, Medicine or Physiology 1991–1995* N Ringertz, ed, pp 314–326 (World Scientific, Singapore, 1997).

The classical ideas about pattern formation in non-equilibrium systems were presented by Turing (1952), who was aiming specifically at an understanding of embryonic development. Modern views are given by Cross & Hohenberg (1993) and by Cross & Greenside (2009).

Cross & Greenside 2009: *Pattern Formation and Dynamics in Nonequilibrium Systems* M Cross & H Greenside (Cambridge University Press, Cambridge 2009).

Cross & Hohenberg 1993: Pattern formation outside of equilibrium. MC Cross & PC Hohenberg, *Revs Mod Phys* **65**, 851–1112 (1993).

Turing 1952: The chemical basis of morphogenesis. AM Turing, *Phil Trans R Soc Lond B* **237**, 33–72 (1952).

The general idea that cells know their position, and hence their fate, in an embryo by responding to the concentration of some special “morphogen” molecule is very old, and it didn’t take too

long before people started to think about the role of diffusion in establishing morphogen gradients. Some milestones are Wolpert’s discussion of positional information (Wolpert 1969), and Crick’s surprisingly influential discussion of diffusion (Crick 1970). The transcription factor bicoid, in the *Drosophila* embryo, provides a very clear example of these ideas (Driever & Nüsslein-Vollhard 1988a,b; Ephrussi & St Johnston 2004). **I am embarrassed not to know who first wrote down the simple model for Bcd profiles, and I should check!**

Crick 1970: Diffusion in embryogenesis. F Crick, *Nature* **225**, 420–422 (1970).

Driever & Nüsslein-Vollhard 1988a: A gradient of Bicoid protein in *Drosophila* embryos. W Driever & C Nüsslein-Vollhard, *Cell* **54**, 83–93 (1988).

Driever & Nüsslein-Vollhard 1988b: The Bicoid protein determines position in the *Drosophila* embryo in a concentration-dependent manner. W Driever & C Nüsslein-Vollhard, *Cell* **54**, 95–104 (1988).

Ephrussi & St Johnston 2004: Seeing is believing: The bicoid morphogen gradient matures. A Ephrussi & D St Johnston, *Cell* **116**, 143–152 (2004).

Wolpert 1969: Positional information and the spatial pattern of cellular differentiation. L Wolpert, *J Theor Biol* **25**, 1–47 (1969).

Houchmandzadeh et al (2002) drew attention to the problem of variability in morphogen gradients, and their suggestion that the emergence of reproducible patterns was an example of robustness in biochemical networks attracted considerable attention. Among the models that emerged in an attempt to flesh out the idea of robustness, some make specific use of gradients from the two ends of the embryo to compensate for global parameter variations and allow for scaling with the size of the egg (Houchmandzadeh et al 2005, McHale et al 2006), while others use nonlinearities in degradation reactions (Eldar et al 2002) or in the transport process (Bollenbach et al 2005) to generate spatial profiles that are robust against variations in source strength. Although much of this discussion focuses on early events in embryonic development, there is also the idea that the final patterns of gene expression, which are more closely tied to cell fate, should be robust steady states of the relevant biochemical networks (von Dassow et al 2000). Even earlier work emphasized the similarity of these networks to neural nets, with stable patterns being analogous to stored memories (Mjolsness et al 1991), and one can see this as a modern formulation of the ideas of “canalization” (Waddington 1942). **Most recent work from Sig-gia & Carlson. Have to see what gets said about size control, but certainly will cite Shraiman (2005).**

Bollenbach et al 2005: Robust formation of morphogen gradients. T Bollenbach, K Kruse, P Pantazis, M González-Gaitán & F Jülicher, *Phys Rev Lett* **94**, 018103 (2005).

von Dassow et al 2000: The segment polarity network is a robust developmental module. G von Dassow, E Meir, EM Munro & GM Odell, *Nature* **406**, 188–192 (2000).

Eldar et al 2002: Robustness of the BMP morphogen gradient in *Drosophila* embryonic development. A Eldar, R Dorfman, D Weiss, H Ashe, B-Z Shilo & N Barkai, *Nature* **419**, 304–308 (2002).

Houchmandzadeh et al 2002: Establishment of developmental precision and proportions in the early *Drosophila* embryo. B Houchmandzadeh, E Wieschaus & S Leibler, *Nature* **415**, 798–802 (2002).

Houchmandzadeh et al 2005: Precise domain specification in the developing *Drosophila* embryo. B Houchmandzadeh, E Wieschaus & S Leibler, *Phys Rev E* **72**, 061920 (2005).

McHale et al 2006: Embryonic pattern scaling achieved by oppositely directed morphogen gradients. P McHale, W-J Rappel & H Levine, *Phys Biol* **3**, 107–120 (2006).

Mjolsness et al 1991: A connectionist model of development. E Mjolsness, DH Sharp & J Reintz, *J Theor Biol* **152**, 429–453 (1991).

Shraiman 2005: Mechanical feedback as a possible regulator of tissue growth. BI Shraiman, *Proc Nat'l Acad Sci (USA)* **102**, 3318–3323 (2005).

Waddington 1942: Canalization of development and the inheritance of acquired characters. CH Waddington, *Nature* **150**, 563–565 (1942).

For measurements on the reproducibility of the early events in the fly embryo, see Dubuis et al (2011). **Are there classical references?** As noted above, and in Section 2.3, the overall precision and reproducibility of pattern formation in the fruit fly embryo is equivalent to $\sim 10\%$ accuracy in the concentration of Bicoid. Although this might not be how things actually work, it does suggest a standard for making measurements of the Bicoid concentration (and, perhaps, for other morphogens as well). For a recent discussion of the state of the art in these experiments, see Dubuis et al (2010). The measurements on reproducibility of the Bcd profiles shown in Fig 121 are from Gregor et al (2007b), cited in Section **[** General decision—is it ok to give references more than once in different sections?]**. Experiments on the scaling of Bcd profile across species include Gregor et al (2005) and Gregor et al (2008).

Dubuis et al 2010: Quantifying the Bicoid morphogen gradient in living fly embryos. J Dubuis, AH Morrison, M Scheeler & T Gregor, arXiv:1003.5572 [q-bio.QM] (2010).

Dubuis et al 2011: Positional information, in bits. JO Dubuis, G Tkačik, W Bialek, EF Wieschaus & T Gregor, in preparation (2011).

Gregor et al 2005: Diffusion and scaling during early embryonic pattern formation. T Gregor, W Bialek, DW Tank, RR de Ruyter van Steveninck, DW Tank & EF Wieschaus, *Proc Nat'l Acad Sci (USA)* **102**, 18403–18407 (2005).

Gregor et al 2008: Shape and function of the Bicoid morphogen gradient in dipteran species with different sized embryos. T Gregor, AP McGregor & EF Wieschaus, *Dev Biol* **316**, 350–358 (2008).

D. Long time scales in neural networks

The basic time scales of electrical dynamics in neurons are measured in milliseconds, yet the time scales of our mental experience are much longer. From the fraction of a second that we need to integrate sounds as we identify words or phrases, to the minutes of memory for a phone number, to the decades over which our recollections of childhood experiences can stretch, the brain has access to time scales far beyond those describing the elementary events of action potential generation and synaptic transmission. If we write a set of dynamical equations, and the time scales which emerge to describe the whole system are much longer than the time scales which appear as parameters in the equations, then something special has happened. How does this work in the brain? How does the system insure that this seemingly special separation of time scales occurs robustly?

One possible solution to the wide range of relevant time scales is to invoke a correspondingly wide range of mechanisms, and surely this is part of the right answer. Thus, it seems unlikely that memories of things long past are stored as continuing patterns of electrical activity in the brain, which somehow last for $\sim 10^{10} \times$ longer than their natural time scale, and are always present to be examined as we reminisce. On the other hand, for working memory—holding the words of a sentence in our minds, or doing mental arithmetic—the time scales involved seem at once long compared with natural time scales for electrical activity, yet too short to engage biochemical mechanisms, such as the regulation of gene expression, which could have more stable, semi-permanent effects.

In fact, we know a whole class of examples in which long time scales emerge naturally. When a ball rolls down a hill, the time scale of the rolling may be short, but once at the bottom the ball can stay there (more or less) forever. So, perhaps we can arrange for the dynamics of neurons in an interconnected network to be like the motion of a particle on a (multidimensional) landscape, with nice deep valleys corresponding to patterns of activity that can persist for a long time once the system find itself in the right neighborhood. In two hugely influential papers in the early 1980s, Hopfield showed how to do exactly this.

A typical neuron in the brain receives inputs from many other neurons **[need to see where we've had a chance to talk about axons, dendrites, synapses .. should be before this!]**; in the cortex 'many' is several thousand, and in the extreme case of the cerebellum 'many' actually means $\sim 10^5$. Conversely, although each cell has only one axon along which its output action potentials are sent, this axon can branch to contact thousands of




FIG. 123 A schematic network of neurons, focusing on one cell i that receives inputs from many other cells $j = 1, 4, 8, \dots$.

other cells. Let's focus on one cell i , which receives inputs from many other cells j , as in Fig 156. Schematically, we can imagine that each cell is either active or inactive, on or off, and hence the state of one cell can be represented by a binary variable $\sigma_i = \pm 1$; for the moment we will leave this as schematic, and not try to interpret σ_i too closely in terms of action potentials or membrane voltage. In the simplest view, each cell j sends its output to cell i , and as these inputs are collected from the synapses, they are summed with some weights W_{ij} which we can think of as the "strengths" of the synapses from cell j to cell i . Having summed its inputs, cell i must then decide whether to be on or off, comparing the total input to a threshold θ_i . These words are equivalent to saying that the state of cell i is set according to the equation

$$\sigma_i \rightarrow \operatorname{sgn} \left[\sum_j W_{ij} \sigma_j - \theta_i \right]. \quad (576)$$

Models of this flavor go back at least to the 1940s, when McCulloch and Pitts explored the idea that the on/off states of neurons could implement a kind of logical calculus. Precisely because they can perform such operations, these sorts of discrete dynamics can be almost arbitrarily complicated. Thus, in general, it's hard to say anything about the dynamics generated by Eq (576).

Suppose, however, that if neuron j synapses onto neuron i with strength W_{ij} , then neuron i synapses onto neuron j with the same strength, so that the matrix of synaptic strengths W_{ij} is symmetric. Then the updating of the state of neuron i in Eq (576) serves to reduce an 'energy' function defined by

$$E = -\frac{1}{2} \sum_{ij} \sigma_i W_{ij} \sigma_j + \sum_i \theta_i \sigma_i. \quad (577)$$

Indeed, we recognize Eq (576) as being the dynamics of a zero temperature Monte Carlo simulation of an Ising model with energy defined by Eq (577). Now, we can make progress.

Problem 120: Energy in the Hopfield model. Show explicitly that the dynamics in Eq (576) serves to decrease the energy function in Eq (577).

If we can map the dynamics of a neural network onto the Ising model, then we can bring an enormous amount of our intuition (and mathematical tools) from statistical mechanics. We know that, since the dynamics we have defined are at zero temperature—we are neglecting, for the moment, any noise in the neurons or synapses—it

is possible to have collective states of the whole system which are stable forever. The simplest example is with all thresholds equal to zero, and all synaptic strengths equal and positive. Then the energy function becomes

$$E = -\frac{W}{2} \sum_{ij} \sigma_i \sigma_j = -\frac{W}{2} \left(\sum_i \sigma_i \right)^2. \quad (578)$$

This is the mean-field ferromagnet. In this model there are two stable ground states—all neurons 'on' ($\sigma_i = +1$ for all i) and all neurons 'off' ($\sigma_i = -1$ for all i). Two states aren't many, and these states seem especially odd, but maybe we are on the right track.

If instead of making all the W_{ij} equal, we choose them at random, then the Ising model we have constructed is the mean-field or Sherrington-Kirkpatrick spin glass. We know that this system has many locally stable states, with an energy landscape that has valleys within valleys, as discussed in Section III.A. This is probably too much, since the structure of these exponentially large number of states depends very sensitively on the precise form of the couplings W_{ij} . More generally, since we only have $\sim N^2$ parameters at our disposal when we adjust the W_{ij} , it is difficult to imagine how we could 'program' the network to store exponentially many independent patterns.

To find a compromise between the ferromagnet and the spin glass, we recall a trick from the history of models for magnetism. Suppose that

$$W_{ij} = W \xi_i \xi_j, \quad (579)$$

where $\vec{\xi}$ is an arbitrary binary vector, $\xi_i = \pm 1$, and for simplicity let the thresholds $\theta_i = 0$. Then the energy becomes

$$\begin{aligned} E &= -\frac{1}{2} \sum_{ij} \sigma_i W_{ij} \sigma_j \\ &= -\frac{W}{2} \sum_{ij} \sigma_i \xi_i \xi_j \sigma_j \end{aligned} \quad (580)$$

$$= -\frac{W}{2} \sum_{ij} (\xi_i \sigma_i) (\xi_j \sigma_j) \quad (581)$$

$$= -\frac{W}{2} \sum_{ij} \tilde{\sigma}_i \tilde{\sigma}_j, \quad (582)$$

where $\tilde{\sigma}_i = \xi_i \sigma_i$ is again a binary variable, $\tilde{\sigma}_i = \pm 1$. The transformation $\sigma_i \rightarrow \tilde{\sigma}_i$ is a discrete gauge transformation, so we see that the model with weights in Eq (579) is gauge equivalent to a ferromagnet. Rather than the stable states of the system being $\sigma_i = +1$ for all i and $\sigma_i = -1$ for all i , the stable states are $\sigma_i = +\xi_i$ and $\sigma_i = -\xi_i$. Importantly, this construction can be generalized.

Rather than Eq (579), let us imagine that

$$W_{ij} = W \left(\xi_i^{(1)} \xi_j^{(1)} + \xi_i^{(2)} \xi_j^{(2)} \right). \quad (583)$$

Now we have

$$E = -\frac{1}{2} \sum_{ij} \sigma_i W_{ij} \sigma_j$$

$$= -\frac{W}{2} \left[\sum_{ij} \sigma_i \xi_i^{(1)} \xi_j^{(1)} \sigma_j \right] - \frac{W}{2} \left[\sum_{ij} \sigma_i \xi_i^{(2)} \xi_j^{(2)} \sigma_j \right]$$
(584)

$$= -\frac{W}{2} \left[\left(\vec{\xi}^{(1)} \cdot \vec{\sigma} \right)^2 + \left(\vec{\xi}^{(2)} \cdot \vec{\sigma} \right)^2 \right].$$
(585)

Clearly the energy will be low if the pattern of neural activity $\vec{\sigma}$ is parallel to the vector $\vec{\xi}^{(1)}$ or to the vector $\vec{\xi}^{(2)}$. But in a high dimensional space, two randomly chosen vectors are, with high probability, nearly orthogonal. This means that the two terms in the Hamiltonian can't both be important at once. Thus, the energy function will have a minimum near $\vec{\sigma} = \vec{\xi}^{(1)}$ and a separate minimum near $\vec{\sigma} = \vec{\xi}^{(2)}$, as well as the flipped versions of these states, $\vec{\sigma} = -\vec{\xi}^{(1)}$ and $\vec{\sigma} = -\vec{\xi}^{(2)}$.

Problem 121: Random vectors in high dimensions.

Consider random binary vectors \vec{v} in an N -dimensional space: $\vec{v} \equiv \{v_1, v_2, \dots, v_N\}$, where each $v_i = \pm 1$ is chosen independently and at random. The angle ϕ between two such vectors is defined in the usual way by normalizing the dot product,

$$\cos \phi \equiv \frac{1}{N} \vec{v}^{(1)} \cdot \vec{v}^{(2)}. \quad (586)$$

Before calculating anything, explain why, if $\vec{v}^{(1)}$ and $\vec{v}^{(2)}$ are chosen independently, it must be that $\langle \cos \phi \rangle = 0$. Calculate the variance $\langle \cos^2 \phi \rangle$ to show that the typical values of $|\cos \phi| \sim 1/\sqrt{N}$, which vanishes as $N \rightarrow \infty$. Can you use the central limit theorem to say something about the whole probability distribution $P(\cos \phi)$ in this limit? Show that the distribution can be written exactly as

$$P(z = \cos \phi) = \int \frac{dk}{2\pi} e^{-ikz} [\cos(k/N)]^N. \quad (587)$$

Connect this result to the predictions of the central limit theorem. Develop a saddle point approximation so that you can calculate, at large N , $P(z)$ for values of $|z| \gg 1/\sqrt{N}$. Verify your approximations with a simulation.

The key idea now is to go further, with not just two patterns but many, writing the weights as

$$W_{ij} = W \sum_{\mu=1}^K \xi_i^{(\mu)} \xi_j^{(\mu)}. \quad (588)$$

Then the energy becomes

$$E = -\frac{W}{2} \sum_{ij} \sigma_i \left[\sum_{\mu=1}^K \xi_i^{(\mu)} \xi_j^{(\mu)} \right] \sigma_j = -\frac{W}{2} \sum_{\mu=1}^K \left(\vec{\xi}^{(\mu)} \cdot \vec{\sigma} \right)^2. \quad (589)$$

Certainly if $K \ll N$ our intuition from the case of two patterns should carry over, since almost all of the vectors $\vec{\xi}^{(\mu)}$ will be nearly orthogonal, and we should find that the energy function has $2K$ minima, near the vectors $\pm \vec{\xi}^{(\mu)}$. At some value of K this must stop being true; indeed if we let K itself become large we must get back to the spin glass model in which there are many locally stable states, but they don't have any connection to the patterns $\vec{\xi}^{(\mu)}$ that we have 'programmed' into the system. In his original work on this model, Hopfield gave rough arguments to suggest that this transition from ordered to disordered behavior occurs at roughly $K \sim 0.15N$, so that it should be possible to have a number of states which is proportional to the number of neurons, and he verified this in simulations with $N = 100$ [should break this off as a paragraph and give the argument, rather than pointing].

Problem 122: Simulating the Hopfield model. Given a matrix W_{ij} it is straightforward to simulate the dynamics of the Hopfield model, as defined by Eq (576); try the simplest case, with $\theta_i = 0$. To run the simulation, you can go through these steps:

1. Start a collection of N spins in some randomly chosen state.
2. Choose one spin i at random.
3. Set $\sigma_i = \text{sgn} \left[\sum_j W_{ij} \sigma_j \right]$.
4. Choose another spin and repeat the update, again and again

Produce a series of simulations to convince yourself that, with W_{ij} chosen as in Eq (588) and a small value of K , the dynamics always stop in the neighborhood of one of the vectors $\vec{\xi}^{(\mu)}$ that you have used in sculpting the energy landscape. Explore what happens as K becomes larger. If you jump to $K \sim N/2$, can you see the emergence of more random stopping points for the dynamics? Perhaps even if you start at one of the vectors $\vec{\xi}^{(\mu)}$, the interference from the other vectors destabilizes this state? If the dynamics stops at a state $\vec{\sigma}_s$, define an order parameter by finding the nearest vector $\vec{\xi}^{(\mu)}$, and measuring the normalized dot product,

$$m_s = \max_{\mu} \left| \vec{\xi}^{(\mu)} \cdot \vec{\sigma}_s \right|. \quad (590)$$

From many random starting points, what is the mean value of m_s as a function of K and N ? As N gets larger, do you see the emergence of a 'thermodynamic limit,' where the (intensive) order parameter $\langle m_s \rangle$ depends only on the ratio K/N ? Are there signs of a phase transition at some critical value of K/N ?

The idea that the dynamics of neural networks could be mapped onto the Ising model immediately captured the imagination of the physics community. But before exercising ourselves in this direction, let's think about how much progress we have made toward solving our original problem. The Hopfield model shows how the dynamics of a neural network can correspond to 'downhill' motion on an energy landscape, much like a ball rolling down a

hill. Thus, the system as a whole has collective, macroscopic states which will persist for times arbitrarily long compared with the basic time scales of the system, the time scales on which the individual neurons update their microscopic states according to Eq (576). Importantly, there are not just a few of these stable states, but many, in proportion to the number of neurons. Unlike the case of the ball coming to a stop at the bottom of the hill, the stability of these states is the result of activity, each neuron receiving continuous input from other neurons in the network; in effect the stable states are patterns of electrical activity which can reinforce themselves as they propagate through the network, embodying old ideas about the ‘reverberation’ of activity patterns through the extensive feedback loops found in the brain.

It is tempting to think of the stable patterns of activity, $\vec{\sigma} \approx \vec{\xi}^{(\mu)}$ as *being* memories. When we set the synaptic connection matrix to the form shown in Eq (588), we “store” the memories, and as the dynamics settles into one of its locally stable states, one of these memories is “recalled.” Each of the stored memories has a large basin of attraction, so the network will recall the memory given only a relatively weak “hint” that the memory is somewhere in the neighborhood of the current state. I use quotation marks extensively here to highlight the fact that we are sliding from properties of the equations into the everyday language that we use in describing our internal mental experiences, and this is dangerous. But, of course, it is also great fun.

A crucial property of the model is that a particular memory—e.g., $\mu = 42$ —is not stored in any particular place. There is no single neuron or synapse that has responsibility for remembering this single recallable item. Instead, the memory is distributed over essentially all of the elements in the system. Correspondingly, if we eliminate one neuron or one synapse, there is no catastrophic loss of one memory, but at worst a gentle degradation of all the memories; in the limit $K \ll N$ we might even imagine that, as $N \rightarrow \infty$ deletion of anything less than a finite fraction of cells or synapses would have a vanishingly small effect. This ‘fault tolerance’ is a highly attractive property.

Problem 123: Fault tolerance. Develop a small simulation to illustrate the idea of fault tolerance in the Hopfield model.

One might worry that all of this depends upon a very particular form of the synaptic weight matrix, Eq (588). But this form is both natural and, perhaps surprisingly, well connected to experiment. Suppose that the current state of activity in the network, $\vec{\sigma}(t)$ represents something

that we would like to store and be able to recall later. If every synaptic strength is changed by the rule

$$W_{ij} \rightarrow W_{ij} + W \sigma_i(t) \sigma_j(t), \quad (591)$$

then, assuming that we have not already tried to store too many patterns in the network, the current state $\vec{\sigma}(t)$ will act as one more pattern that can be recalled, one more stable state in the energy landscape—the network will have “learned” the state $\vec{\sigma}(t)$. Importantly, the change in strength of the synapse from neuron i to neuron j depends only on the states of neurons i and j . Thus, although the memory is distributed throughout the network, the rule for storing the memory is completely local.

The rule for modification of synaptic strengths in Eq (591), sometimes called a “learning rule,” means that, over time, the strength of the synapse from neuron j to neuron i will be proportional to the correlation between the activities of these two cells. Learning based on correlations is an idea that goes back at least to Hebb in the 1940s, although there are clear precursors in the writing of William James fifty years earlier. Both James and Hebb were making an intuitive leap between the macroscopic phenomena of human and animal learning and what they imagined could be the underlying neural mechanisms. Although their words admit some breadth of interpretation, to a remarkable extent they were right, and many synapses are found to exhibit “Hebbian plasticity.”

At this point we should say something about the experiments which demonstrate Hebbian plasticity at real synapses. Should get as far as explaining that there is a new issue of time scale separation, since the memory trace should be written quickly (so that the relevant biochemical mechanisms must switch quickly) but then be stable for long times, despite the fact that all the molecules get replaced fairly often. Models for this bring us back to the question of stability against noise in biochemical networks, which is something that should have been covered, in part in Section II.B. There is a lot that one could say here (one could make a nice course about synaptic plasticity alone), so careful selection is required.

What is the evidence that something like the Hopfield model is actually a correct description of real neural networks? The essence of the model, shorn of the analogies to magnetism, is that a recalled memory is a stable state of neural activity, one which persists in the absence of external stimuli by reverberating in the network. Persistent activity of neurons has been observed. The canonical example occurs when an animal has to remember a sensory stimulus for a brief time (a few second to a minute) in order to compare it with another image or more simply because an immediate response would be impossible. In the period between the stimulus and the cue for the response, where the subject has to remember what has been seen or heard, these neurons continue to generate action potentials at a rate very different from the ‘resting’ rate

before exposure to the initial stimulus, as shown in Fig 124. Although the behavior of each cell is different in detail, in many cases the activity during this ‘delay period’ is steady, as if the system were simply locked into a new state, but the state into which the system falls is different depending on the image which is being remembered. Persistent activity is not just a feature of our cortex, but appears also in many other systems, from the primate spinal cord to the goldfish brainstem. [Probably need more here: demonstrate that persistent activity varies in relation to the triggering inputs, in some cases is continuously graded, etc..]

One would like to demonstrate directly that the persistent activity of individual neurons during the delay period really reflects a collective state of the network. This is not so easy to do. Need to decide how far to go here—are there good experiments in cortex looking at synaptic inputs? Maybe say that this is an important reason to look for simpler examples ... Also want at least to point toward Amit’s analysis of the Miyashita correlations, where the persistent patterns of activity have a trace of the sequence in which images were presented during learning.

At this point it would be nice to say a little about the more sophisticated analysis of the Hopfield model using replicas. The goal is to calculate the ‘capacity,’ that is the maximum number of patterns K that can be stored




FIG. 124 The activity of a single neural in primate prefrontal cortex during short-term memory, from Fuster & Alexander (1971). In these experiments a rhesus monkey is trained to open one of two doors when he receives a cue that they are unlocked (response cues). Some time before this, the subject has been allowed to see which of the doors has a piece of apple behind it (stimulus cue). This neuron seems to be active during the delay period, and this persistent activity plausibly is part of the memory that the subject hold. These data record the results of five such experiments, where the vertical lines mark the times of spikes, and the arrows mark the times of the cues, as labelled.

and successfully retrieved. This can be formulated as a problem in the statistical mechanics of disordered systems. I am not sure how much technical force is needed here (or in the discussion of protein above). Advice is welcome!

There is a very different way of connecting the Hopfield model to experiment. Imagine that we divide time into small bins of duration $\Delta\tau$. If $\Delta\tau$ is sufficiently small, then each neuron either generate an action potential in this bin, or it does not, so that the neural response is naturally binary: $\sigma_i = +1$ for a spike, $\sigma_i = -1$ for silence. For a large network it is impossible to ‘measure’ the probability distribution of all the network states, $P(\vec{\sigma})$. But even recording from neurons one by one it is possible to measure the mean rate at which each cell generates spikes, which is equivalent to the expectation value $\langle\sigma_i\rangle$, and it is becoming increasingly common to record at least from pairs of cells, which makes it possible to estimate the correlations $C_{ij} \equiv \langle\sigma_i\sigma_j\rangle - \langle\sigma_i\rangle\langle\sigma_j\rangle$. One could ask, as a purely practical question, what do these measurements tell us about the full distribution $P(\vec{\sigma})$? In general, of course, there are infinitely many distributions (over the 2^N states) that are consistent with these $N(N+1)/2$ measurements. Out of all these possible distributions, there is one which reproduces the measurements but otherwise describes a system which is as random or unstructured as possible, and this is the maximum entropy distribution, as we discussed in Section III.A; see also Appendix A.8.

We recall that the maximum entropy distribution consistent with a certain mean energy for a system is the Boltzmann distribution. This construction generalizes. Suppose that we are looking for the probability distribution $P(\vec{\sigma})$, and we know the expectation values of some functions on the state, $\langle f_\mu(\vec{\sigma}) \rangle = \bar{f}_\mu$. Then to maximize the entropy of the distribution subject to these constraints, we use Lagrange multipliers as usual. Thus, our problem is to maximize [again, let’s be careful about how this is done here vs. earlier vs. Appendix A.8]

$$\mathcal{F} = - \sum_{\vec{\sigma}} P(\vec{\sigma}) \ln P(\vec{\sigma}) - \sum_{\mu} \lambda_{\mu} \left[\sum_{\vec{\sigma}} P(\vec{\sigma}) f_{\mu}(\vec{\sigma}) - \bar{f}_{\mu} \right] - \Lambda \left[\sum_{\vec{\sigma}} P(\vec{\sigma}) - 1 \right], \quad (592)$$

where the last term fixes the normalization of the distribution. Following through the steps, the optimum is defined by

$$0 = \frac{\delta \mathcal{F}}{\delta P(\vec{\sigma})} = - [\ln P(\vec{\sigma}) + 1] - \sum_{\mu} \lambda_{\mu} f_{\mu}(\vec{\sigma}) - \Lambda \quad (593)$$

$$\ln P(\vec{\sigma}) = - \sum_{\mu} \lambda_{\mu} f_{\mu}(\vec{\sigma}) - (\Lambda + 1) \quad (594)$$

$$P(\vec{\sigma}) = \frac{1}{Z} \exp \left[- \sum_{\mu} \lambda_{\mu} f_{\mu}(\vec{\sigma}) \right], \quad (595)$$

where the partition function $Z = e^{-(\Lambda+1)}$, or, fixing normalization,

$$Z(\{\lambda_\mu\}) = \sum_{\vec{\sigma}} \exp \left[- \sum_{\mu} \lambda_{\mu} f_{\mu}(\vec{\sigma}) \right]. \quad (596)$$

The multipliers λ_{μ} are determined by matching the expectation values in the distribution to those observed experimentally. We recall the usual identity

$$\langle f_{\nu}(\vec{\sigma}) \rangle = - \frac{\partial \ln Z(\{\lambda_{\mu}\})}{\partial \lambda_{\nu}}, \quad (597)$$

so we have to solve the equations

$$- \frac{\partial \ln Z(\{\lambda_{\mu}\})}{\partial \lambda_{\nu}} = \bar{f}_{\nu} \quad (598)$$

to complete the construction of the model; in general this is a hard task, the inverse of what we usually do in statistical mechanics.

If the expectation values that we measure are $\langle \sigma_i \rangle$ and $\langle \sigma_i \sigma_j \rangle$, then the corresponding maximum entropy distribution can be written as

$$P(\vec{\sigma}) = \frac{1}{Z} \exp \left[\sum_{i=1}^M h_i \sigma_i + \frac{1}{2} \sum_{i \neq j}^N J_{ij} \sigma_i \sigma_j \right], \quad (599)$$

where the ‘magnetic fields’ $\{h_i\}$ and the ‘exchange couplings’ $\{J_{ij}\}$ have to be set to reproduce the measured values of $\{\langle \sigma_i \rangle\}$ and $\{\langle C_{ij} \rangle\}$. This of course is an Ising model with pairwise interactions among the spins. What is crucial is that this model emerges here not through hypotheses about the network dynamics, but rather as the least structured model that is consistent with the measured expectation values. The mapping to the Ising model is a mathematical equivalence, not an analogy, and the details of the model are specified by the data.

The emergence of the Ising model is an attractive aspect of the maximum entropy construction. But, there is no obvious reason why real biological networks should have this maximum entropy property. Indeed, one might guess that there are complicated, higher order correlations which are important for the function of the network, and these will be missed by a maximum entropy model built only from pairwise correlations. It thus came as a surprise when it was found that these models really do provide an accurate description of the full correlation structure in the vertebrate retina as it responds to naturalistic stimuli. This has led to considerable interest in the use of these models more generally for the description of real neural networks; for details, see Appendix A.8.

Problem 124: Maximum entropy model for a simple neural network. Imagine that we record from N neurons and we

find that all of them have the same mean rate of spiking, \bar{r} . Further, if we look at any pair of neurons, the probability of both spiking in the same small window of duration $\Delta\tau$ is $p_c = (\bar{r}\Delta\tau)^2(1 + \epsilon)$. We want to describe this network as above, with Ising variables $\sigma_i = +1$ for spiking and $\sigma_i = -1$ for silence.

(a.) Show that

$$\langle \sigma_i \rangle = -1 + \bar{r}\Delta\tau \quad (600)$$

$$C_{ij} \equiv \langle \sigma_i \sigma_j \rangle - \langle \sigma_i \rangle \langle \sigma_j \rangle = 4\epsilon(\bar{r}\Delta\tau)^2. \quad (601)$$

(b.) Since all neurons and pairs are equivalent, the maximum entropy model consistent with pairwise correlations has the simpler form,

$$P(\vec{\sigma}) = \frac{1}{Z} \exp \left[h \sum_{i=1}^M \sigma_i + \frac{J}{2} \sum_{i \neq j}^N \sigma_i \sigma_j \right], \quad (602)$$

which is just the mean field ferromagnet (assuming that J is positive). If N is large, one might expect that there is a ‘thermodynamic limit’ in which quantities like energy and entropy become extensive, proportional to N . Show that this requires scaling of the coupling, $J = J_0/N$. With this scaling, derive the relationship between the derivatives of $\ln Z$ and the expectation values $\langle \sigma_i \rangle$ and C_{ij} .

(c.) Some of you will be very familiar with the substitution tricks that we’re about to use, others less so. To be sure, let me take you through the steps. We notice that the interactions are described by a term

$$\frac{J}{2} \sum_{i \neq j}^N \sigma_i \sigma_j = \frac{J}{2} \sum_{i,j=1}^N \sigma_i \sigma_j - \frac{NJ}{2} = \frac{J}{2} \left(\sum_{i=1}^N \sigma_i \right)^2 - \frac{NJ}{2}. \quad (603)$$

Thus the partition function can be written as

$$Z = \sum_{\vec{\sigma}} \exp \left[h \sum_{i=1}^M \sigma_i + \frac{J}{2} \sum_{i \neq j}^N \sigma_i \sigma_j \right] \quad (604)$$

$$= e^{-NJ/2} \sum_{\vec{\sigma}} \exp \left[h \sum_{i=1}^M \sigma_i \right] \exp \left[\frac{J}{2} \left(\sum_{i=1}^N \sigma_i \right)^2 \right]. \quad (605)$$

Then the key step is to realize that

$$\exp \left[\frac{A}{2} (x)^2 \right] = \int \frac{d\phi}{\sqrt{2\pi A}} \exp \left[-\frac{\phi^2}{2A} + \phi x \right]. \quad (606)$$

Applied to Eq (605) this allows us to write

$$\begin{aligned} Z &= e^{-NJ/2} \sum_{\vec{\sigma}} \exp \left[h \sum_{i=1}^M \sigma_i \right] \exp \left[\frac{J}{2} \left(\sum_{i=1}^N \sigma_i \right)^2 \right] \\ &= e^{-NJ/2} \sum_{\vec{\sigma}} \exp \left[h \sum_{i=1}^M \sigma_i \right] \int \frac{d\phi}{\sqrt{2\pi J}} \exp \left[-\frac{\phi^2}{2J} + \phi \sum_{i=1}^N \sigma_i \right] \end{aligned} \quad (607)$$

$$= e^{-NJ/2} \int \frac{d\phi}{\sqrt{2\pi J}} \exp \left[-\frac{\phi^2}{2J} \right] \sum_{\vec{\sigma}} \exp \left[(h + \phi) \sum_{i=1}^N \sigma_i \right]. \quad (608)$$

Now we see that the spins have decoupled, and you should be able to do the sum over states, $\sum_{\vec{\sigma}}$, inside the integral. Show that, with the scaling from (b.),

$$Z = e^{-NJ/2} \int \frac{d\phi}{\sqrt{2\pi J}} \exp \left[-NF(\phi; h, J_0) \right], \quad (609)$$

where the effective free energy $F(\phi; h, J_0)$ has no explicit N dependence.

(d.) Use steepest descent to approximate Eq (609) at large N . Derive an expression for $\ln Z$ which captures both the leading behavior ($\ln Z \propto N$) and the first two corrections.

(e.) To finish the construction of the model, we have to adjust h and J to match the measured means and pairwise correlations, Eq’s

(600) and (601). Using the scaling required for a thermodynamic limit, is there a prediction for the N dependence of the correlation strength ϵ ? This should bother you— ϵ is a quantity that is *measured* from pairs of cells, and shouldn't really depend on the number of cells in the network. Suppose we measure ϵ among more and more pairs of cells, so we have to describe larger and larger networks. Is it possible to have ϵ small and constant as $N \rightarrow \infty$? What conditions need to be met in order for this to happen?

The Hopfield model provides a scheme for the stabilizing multiple, discrete patterns of activity. But there certainly are situations in which the brain must hold a memory of a continuous variable. This is even less generic than the case of discrete attractors. In order to have a memory of a continuous variable, there must be (at least) a whole line or curve in state space along which the system can stop; if we think it terms of an energy landscape, then there must be one big valley, and the bottom of this valley must be precisely flat along one direction. Implausible as all this sounds, the brain really does hold memories of continuous variables, and it does so even in simple situations.

When you turn your head, cells in the semicircular canals, buried in the same bone as the cochlea, sense the rotational motion; this is called our “vestibular” sense. This angular motion input passes through the brain and drives a motor output which counter-rotates the eyes. This happens automatically, and is called the vestibulo-ocular reflex. You can demonstrate it for yourself by shaking your head from side to side as you read this text. If you are holding the book at arm's length, then in order to read you have to have your fovea—the $\sim 1^\circ$ wide area of highest image quality—focused on the words as you read them. If you move your head from side to side, and don't move your eyes to compensate, the text will blur. In fact, you (hopefully) have no trouble reading and shaking your head at the same time, suggesting that your eyes are being moved to compensate with an accuracy of better than $\sim 1^\circ$. When you are reading, of course, there are visual cues to help guide you eye movements, but it turns out that even if you close your eyes or sit in a dark room, seeing nothing, your eyes still counter-rotate to compensate for your head motions.

Problem 125: Mechanics of the semicircular canals. Give a problem to develop the simple mechanical model of the canal, explaining how one gets velocity sensitivity over a reasonable bandwidth. Use real dimensions of the canal (e.g., in humans) to get numbers out.




FIG. 125 Integration as memory for a continuous variable. Dashed lines show possible velocity signals, and the solid lines show corresponding position signals, obtained by integrating the velocity. After the transient inputs die away, the output of the integrator is stable for all time (a memory) and can take on any real value.

There is a subtlety of the vestibulo-ocular reflex, however. If we relax all the muscles to our eyes, then they rotate to a resting position in which we are looking more or less straight ahead (as defined by where our nose is pointing). Thus, if we turn our head to the right and stop, we need to keep tension on the eye muscles to be sure that they don't drift away from where we were looking before we turned. That is, to fully compensate for rotation of the head we need a signal related to the desired angular *displacement* of the eyes. But the vestibular system is an inertial sensor, driven by angular accelerations; the mechanics of canal turn this into a velocity signal over a wide range of frequencies, but the sensors really have zero response to constant displacements. Thus, the brain needs to take an input related (at best) to head velocity, and generate an output related to head displacement—it has to integrate, where here ‘integrate’ has the literal meaning from calculus, rather than being a qualitative statement about the gathering of multiple signals. Although we don't usually think about it this way, an integrator is a device which, once the input signals die away, has a perfect memory for a continuous variable, as schematized in Fig 125. Although these properties of the integral are obvious mathematically, it is less obvious how to build a network of neurons that implements this mathematics.

Before continuing, it should be noted that the movement of our eyes is not the perfect integral of our head velocity. On a longer time scale, roughly thirty seconds, our eyes do drift back to a resting position if there is no further stimulus. But this time scale is very long compared with the natural time scales of individual neurons, perhaps by a factor of as much as $\sim 10^3$. Could this gap

be closed by an emergent long time scale in the network, resulting from a line or curve of fixed points?

Suppose that the activity of each neuron is described by a coarse-grained continuous variable, such as the rate r at which it generates action potentials. If we inject a current I into the neuron directly, we find that the rate changes, along some curve $r(I)$. Each spike arriving at a synapse onto cell i effectively injects current into that cell, but this current is smoothed by some dynamics which we will summarize by a time scale τ , and the spikes from cell j are weighted by the strength of the synapse W_{ij} . This suggests a simple model,

$$\tau \frac{dI_i}{dt} + I_i = \sum_j W_{ij} r(I_j) + I_i^{\text{ext}}, \quad (610)$$

where I_i^{ext} represents currents injected from outside the network, including from sensory inputs. Typical examples of the response function $g(I)$ are sigmoids, threshold linear relations, and other monotonic functions. [\[add figure to show some examples of \$g\(I\)\$?\]](#)

What would it mean for the dynamics of Eq (610) to be an integrator? At the very least, the dynamics has to look like an integrator in its linear response to inputs, so let's see how this is possible. Assume that in the absence of inputs, there is some steady state at which $I_i = I_i^*$. Then if we linearize around this, writing $I_i = I_i^* + u_i$, we have

$$\tau \frac{du_i}{dt} + u_i = \sum_j W_{ij} r'(I_j^*) u_j + I_i^{\text{ext}}. \quad (611)$$

As always with linear problems, we want to change coordinates so that matrices become diagonal. If we denote quantities in this new coordinate system by tildes, then we will have

$$\tau \frac{d\tilde{u}_n}{dt} + \tilde{u}_n = \Lambda_n \tilde{u}_n + \tilde{I}_n^{\text{ext}}, \quad (612)$$

where the eigenvalues are defined as solutions to

$$\sum_j W_{ij} r'(I_j^*) \psi_j^{(n)} = \Lambda_n \psi_j^{(n)}. \quad (613)$$

If one of the $\Lambda_n \rightarrow 1$, then along this direction we have simply

$$\tau \frac{d\tilde{u}_n}{dt} = +\tilde{I}_n^{\text{ext}}, \quad (614)$$

$$\Rightarrow \tilde{u}_n(t) = \tilde{u}_n(0) + \frac{1}{\tau} \int_0^t dt' \tilde{I}_n^{\text{ext}}(t'), \quad (615)$$

so that \tilde{u}_n is the time integral of its inputs. Thus, being an integrator means arranging the matrix of synaptic strength so that it (in appropriate units) has a unit eigenvalue, which means that (at least in this one mode) the signals which are being received from other cells in the network perfectly balance the decay processes

within each cell. This of course is a critical point in the dynamics—if the eigenvalue is larger than one, the dynamics become unstable, if it is less than one it is stable but an imperfect integrator. Only at the critical point is true integration achieved. If we are within ϵ of the critical point, the system will hold a memory for $\sim \tau/\epsilon$, so if we really need to span three orders of magnitude (or even two), then the adjustment to the critical point must be quite precise.

The language of eigenvalues and critical points makes precise our initial intuition that there is something highly non-generic about memory for a continuous variable. Most valleys have a single lowest point, and balls keep rolling downhill until they find it. Only at the critical point is there one perfectly neutral direction in the valley, along which the ball feels no force.

Problem 126: Details of the line attractor. [\[go through Seung \(1996\) to look for good questions about the linear algebra of the model\]](#)

The fact that the position of our eyes is the integral of the velocity signals from our semicircular canals, and that there is (apparently) a continuum of stable points where our eyes can sit, means that something like this description in terms of line attractors must be true for the system as a whole. Indeed this is more general: the fact that we (and other animals) can stabilize a continuously variable set of postures means that the combined dynamics of our limbs, muscles, sensors and brain must have a line or manifold of attractors. It is more challenging to point to a particular part of the system—e.g., a particular sub-network of neurons in one part of the brain—and claim that the dynamics of this subsystem must have a line or attractors.

Seeing a model which explains things but only for particular choices of parameters makes us uneasy, as in our previous examples in this chapter. But in this case, we know that the relevant parameters—synaptic strengths—are adjustable, because this is how we learn. Also, we know that if we make errors, then under normal conditions (with the lights on) these errors are literally visible as slippage of the image on our retina as we turn our heads. There must be some way to use this error signal to adjust the synaptic weights and tune the network to its critical point. Does this happen?

To test the idea that the brain tunes the dynamics of the integrator circuit to its critical point, Major et al did a seemingly simple but beautiful experiment using goldfish, which also exhibit oculomotor integration. Essentially they built a planetarium for the goldfish, and then




FIG. 126 Schematic of the “planetarium experiment,” from Major et al (2004a). At left, the dynamics of a leaky (top) or unstable (bottom) integrator are evident as exponential decay or growth of eye position. This can be analyzed by plotting eye velocity vs eye position, revealing a straight line with a sign that indicates stability or instability, and a slope that measures the time constant of the system. At right, the planetarium setup, in which eye movements are monitored and fed back to movements of the surrounding scene.

coupled the rotation of this ‘world’ to their eye movements, as in Fig 126. Under normal conditions, when the eyes move by an angle θ , this is equivalent to the world moving the other way by the same angle. But if we give an additional rotation, we can create a situation in which the world slips on the retina even when the integrator network is set correctly. If the system in fact continuously uses slip signals to tune the system, this will drive a mistuning, either toward stability or instability. If we remove the feedback, we should then see that the fish can no longer stabilize its gaze, with the eyes either quickly relaxing to their resting position or exploding wildly away from rest, needing correction by frequent saccades.

The quick summary is that all of what we expect to see is observed experimentally, as summarized in Fig 127. Importantly, one can record from neurons in the relevant circuit and demonstrate that the detuning of the behavioral integration is mirrored by changes in the dynamics of persistent neural firing. While this does not prove that the line attractor scenario is correct, it does show that the long time scale of memory exhibited by the oculomotor integrator is the result of an active tuning process which uses visual feedback as a control signal. In this way, non-generic behavior of the system is learned, robustly.

Need an introduction.




FIG. 127 Results of the planetarium experiment, from Major et al (2004a). At the top (a), control experiments showing the eye trajectories and position vs velocity plots before exposure to the feedback system in the planetarium. Note that the time constant of the system is ~ 20 s. After exposure to feedback which should “teach” the system to be unstable (b) or leaky (c), trajectories and position vs velocity plots show the expected behaviors, with time constants for growth or decay on the order of 1 – 5 s.

Amit 1989: *Modeling Brain Function: The World of Attractor Neural Networks*. DJ Amit (Cambridge University Press, Cambridge, 1989).

Cooper 1973: A possible organization of animal memory and learning. LN Cooper, in *Collective Properties of Physical Systems: Proceedings of Nobel Symposium 24*, B Lundqvist & S Lundqvist, eds, pp 252–264 (Academic Press, New York, 1973).

Hertz et al 1991: *Introduction to the Theory of Neural Computation* J Hertz, A Krogh & RG Palmer (Addison Wesley, Redwood City, 1991).

Hopfield 1982: Neural networks and physical systems with emergent collective computational abilities. JJ Hopfield, *Proc Nat'l Acad Sci (USA)* **79**, 2554–8 (1982).

Hopfield 1984: Neurons with graded response have collective properties like those of two-state neurons. JJ Hopfield, *Proc Nat'l Acad Sci (USA)* **81**, 3088–3092 (1984).

McCulloch & Pitts 1943: A logical calculus of ideas immanent in nervous activity. WS McCulloch & W Pitts, *Bull Math Biophys* **5**, 115–133 (1943).

Now talk about connections to real neurons

Hebb 1949: *The Organization of Behavior: A Neuropsychological Theory*. DO Hebb (Wiley, New York, 1949).

James 1892: *Psychology: The Briefer Course*. W James (Henry Holt and Company, 1892). There is a modern edition from Dover Publications (New York, 2001), based on the 1961 abridged version from Harper and Row (New York, 1961).

Lorente de No 1938: Analysis of the activity of the chains of internuncial neurons. R Lorente de No, *J Neurophysiol* **1**, 207–244 (1938).

Introduce more powerful stat mech approaches

Amit et al 1985: Spin-glass models of neural networks. DJ Amit, H Gutfreund & H Sompolinsky, *Phys Rev A* **32**, 1007–1018 (1985).

Amit et al 1987: Statistical mechanics of neural networks near saturation. DJ Amit, H Gutfreund & H Sompolinsky, *Ann Phys* **173**, 30–67 (1987).

Crisanti et al 1986: Saturation level of the Hopfield model for neural networks. A Crisanti, DJ Amit & H Gutfreund, *Europhys Lett* **2**, 337–341 (1986).

[Pointers to experiments on persistent activity \(need more!\).](#)

Funahashi et al 1989: Mnemonic coding of visual space in the monkey’s dorsolateral prefrontal cortex. S Funahashi, CJ Bruce & PS Goldman–Rakic, *J Neurophysiol* **61**, 331–349 (1989).

Fuster & Alexander 1971: Neuron activity related to short-term memory. JM Fuster & GE Alexander, *Science* **173**, 652–654 (1971).

Prut & Fetz 1999: Primate spinal interneurons show pre-movement instructed delay activity. Y Prut & EE Fetz, *Nature* **401**, 590–594 (1999).

The idea of using maximum entropy models to think about correlations in networks of neurons arose from a very practical problem—if we observe correlations among pairs of neurons, should we be surprised if we observe, for example, three or four neurons generating action potentials simultaneously? For continuous variables, we can separate different orders of correlations quite simply (recall the idea of cumulants in statistics, or “connected diagrams” in field theory). For discrete variables, pairwise correlations imply higher order correlations, even without any further assumptions. One touchstone for this idea is in statistical mechanics—recall that the usual Ising model has only interactions between two spins at a time, but when we coarse grain this model to give the Landau–Ginzburg Hamiltonian, we generate ϕ^4 interaction terms, so that the magnetization ϕ (which is a spatially smoothed version of the original spins) must have nontrivial fourth order correlations [\[should give some standard ref\]](#). Schneidman et al (2003) showed how one could use the maximum entropy construction to generalize the idea of connected correlations to discrete variables. [\[Be careful here .. maybe push more into Appendix A.8?\]](#)

Schneidman et al 2003: Network information and connected correlations. E Schneidman, S Still, MJ Berry II & W Bialek, *Phys Rev Lett* **91**, 238701 (2003).

When we set out to use the maximum entropy method to analyze the responses of real neurons in the vertebrate retina, we expected we would “clean out” the pairwise correlations and uncover the higher order effects which were responsible for the known tendency of many neurons to fire simultaneously (Schnitzer & Meister 2003). The surprising result was that the pairwise Ising model provides a very accurate description of the combinatorial patterns of spiking and silence in ganglion cells of the salamander retina as they respond to natural and artificial movies, and in cortical cell cultures (Schneidman et al 2006). After the initial success in the salamander retina, similarly encouraging results were obtained in the primate retina, under very different stimulus conditions (Shlens et al 2006, 2009), in visual cortex (Ohiorhenuan & Victor 2007, Yu et al 2008), and in networks grown in vitro (Tang et al 2008). Most of these detailed comparisons of theory and experiment were done for groups of $N \sim 10$ neurons, small enough that the full distribution $P_{\text{expt}}(\{\sigma_i\})$ could be sampled experimentally and used to assess the quality of the pairwise maximum entropy model. Attempts to push to larger networks are described by Tkačik et al (2006, 2009) [\[and pointer to Appendix; see how far we can go before finalizing text\]](#).

Ohiorhenuan & Victor 2007: IE Ohiorhenuan & JD Victor, Maximum entropy modeling of multi-neuron firing patterns in V1. Proceedings of 2007 Cosyne conference; <http://cosyne.org>.

Schneidman et al 2006: Weak pairwise correlations imply strongly correlated network states in a neural population.

E Schneidman, MJ Berry II, R Segev & W Bialek, *Nature* **440**, 1007–1012 (2006).

Schnitzer & Meister 2003: Multineuronal firing patterns in the signal from eye to brain. MJ Schnitzer & M Meister, *Neuron* **37**, 499–511 (2003).

Shlens et al 2006: The structure of multi-neuron firing patterns in primate retina. J Shlens, GD Field, JL Gauthier, MI Grivich, D Petrusca, A Sher, AM Litke & EJ Chichilnisky, *J Neurosci* **26**, 8254–66 (2006).

Shlens et al 2009: The structure of large-scale synchronized firing in primate retina. J Shlens, GD Field, JL Gauthier, M Greschner, A Sher, AM Litke & EJ Chichilnisky, *J Neurosci* **29**, 5022–5031 (2009).

Tang et al 2008: A Tang, D Jackson, J Hobbs, W Chen, JL Smith, H Patel, A Prieto, D Petruscam MI Grivich, A Sher, P Hottowy, W Dabrowski, AM Litke & JM Beggs, A maximum entropy model applied to spatial and temporal correlations from cortical networks *in vitro*. *J Neurosci* **28**, 505–518 (2008).

Tkačik et al 2006: Ising models for networks of real neurons. G Tkačik, E Schneidman, MJ Berry II & W Bialek, arXiv:q-bio.NC/0611072 (2006).

Tkačik et al 2009: Spin glass models for networks of real neurons. G Tkačik, E Schneidman, MJ Berry II & W Bialek, arXiv:0912.5409 [q-bio.NC] (2009).

Yu et al 2008: S Yu, D Huang, W Singer & D Nikolic, A small world of neuronal synchrony. *Cereb Cortex* **18**, 2891–2901 (2008).

[\[Refs for oculomotor integrator\]](#)

Major et al 2004a: Plasticity and tuning by visual feedback of the stability of a neural integrator. G Major, R Baker, E Aksay, B Mensh, HS Seung & DW Tank, *Proc Nat’l Acad Sci (USA)* **101**, 7739–7744 (2004).

Major et al 2004b: Plasticity and tuning of the time course of analog persistent firing in a neural integrator. G Major, R Baker, E Aksay, HS Seung & DW Tank, *Proc Nat’l Acad Sci (USA)* **101**, 7745–7750 (2004).

Robinson 1989: Integrating with neurons. DA Robinson, *Ann Rev Neurosci* **12**, 33–45 (1989).

Seung 1996: How the brain keeps the eyes still. HS Seung, *Proc Nat’l Acad Sci (USA)* **93**, 13339–13344 (1996).

E. Perspectives

The exploration of fine tuning vs. robustness in biological systems encourages us to think beyond models for this or that particular system. To ask whether some function requires fine tuning of parameters, we imagine that the system we are looking at is just one member in a class of possible systems. Whatever the answer to our initial questions, this effort at generalization clearly is an important step on the path to a physicist’s view of life.

When we think about individual proteins, generalization is easy—proteins are polymers, and there is a natural class of molecules that can be built from the same

monomers, but with different sequences. When we think about a biochemical or genetic network, with many interacting protein molecules, it seems natural to generalize to a class of networks that has the same topology, but different parameters on each node or link. The ion channels in a single neuron provide an important example of a network of interacting proteins, where the interactions are mediated by the (global) transmembrane voltage and, importantly, experiments on single channel molecules serve to validate the equations describing what happens at each node. Finally, for networks of neurons, the fact that the strengths of synaptic connections are ‘plastic’ makes it natural to think about classes of networks that have the same topology of connections among neurons, but with different strengths. In all of these cases, we can see that the generalization to a class of networks is not just a useful theoretical construct, but also something which has meaning in the life or evolution of the organism.

In the extreme, “robustness” would mean that functional behavior is largely invariant over the whole class of networks. If this really is the case, then we should be able to choose networks at random and have them function. This is essentially the strategy employed by many groups searching for robustness in biochemical networks, and long before this there was a serious exploration of neural networks with randomly chosen strengths of synaptic connections among all the cells, using analytic methods borrowed from the dynamical theory of spin glasses. In the context of neural networks, the model with random connections indeed behaves chaotically, which seems odd, although it has been suggested that in the absence of other inputs this is the right answer—sensory inputs serve to drive the network out of the chaotic phase into an ordered state. For biochemical and genetic networks chaos seems less generic, but to obtain functional behavior without adjusting parameters there is general agreement that the topology of the network must be chosen carefully. There are several open questions here. Why is chaos not more common in large networks of biochemical reactions? What is the boundary between changing parameters (e.g., make the rate on one particular chemical reaction smaller) and changing topology (setting that rate exactly to zero)? To speak precisely about what will be typical of a randomly chosen network, we need a measure on the space of parameters; is there a natural choice of this measure?

In most of the systems we have studied, the randomly chosen parameters do not correspond to functional behavior. Random amino acid sequences don’t fold into functional proteins, randomly chosen numbers of ion channels will not generate the correct rhythms of electrical activity, and while random neural networks may perform some functions, they certainly don’t provide for stable storage and recall of memories. In each of these cases there are mechanisms for tuning or selecting the

functional regions of parameter space. In single neurons, adjusting the numbers of copies of different channels is a form of physiological adaptation, connecting electrical activity, intracellular messengers, and the control of gene expression. In neural networks, the strengths of synapses are adjusted during learning, and for some key processes this learning happens all the time—as perhaps is necessary if the behavior the system is trying to stabilize is very far from typical in the space of possible networks. Finally, for amino acid and DNA sequences, the “adjustment” to functional behavior occurs on evolutionary time scales.⁷⁴ In this context, we can think of adaptation, learning and evolution as different mechanisms for accomplishing the same task, albeit on different time scales.

As we will see in Section IV.D, there is a sophisticated mathematical theory of learning, combining ideas from mathematics, computer science and statistical physics. In particular, in different contexts, this theoretical approach places bounds on what can be learned, and how quickly. If we see adaptation, learning and evolution as different approaches to the same problem, should there be a comparable theoretical framework limiting the speed of evolution, or the effectiveness of adaptation? For evolution there is, in the long run, an obvious external definition of correct functional behavior (successful reproduction), and for learning there are often external signals (as in the case of the oculomotor integrator) that define the goal of the learning process; in adaptation, how do cells “know” the correct behavior that they are trying to stabilize? In the models for regulation of ion channel densities that we discussed in Section III.B, this is (weakly) programmed into the cell by the parameters that define a target calcium concentration; is there a more general definition of when cells are getting things right? Are there, as with learning, limits on how precisely one can get things right if the system needs to adjust quickly?

To return to the opening remarks in this Chapter, we wanted to distinguish between the usual physicist’s mistrust of explanations that rest on fine tuning of parameters, and some specifically biological notions of robustness or evolvability. Part of the motivation for robustness as a biological principle is the intuition that living organisms simply can’t adjust parameters accurately enough to guarantee reliable, reproducible functions. I think this intuition turns out to be wrong—cells can and do exercise precise control over the numbers of molecules that they make, so that the absolute concentrations of relevant molecules *can be* reproducible from cell to cell (or, in the

⁷⁴ It is worth emphasizing that, in the immune system, there is a kind of accelerated evolution within individual organisms, and this serves to select a nontrivial distribution of sequences for the antibody molecules. See the discussion in Section III.A.

discussion of Section III.C, embryo to embryo) with high precision. I emphasize “can be,” because one clearly cannot conclude that all concentrations or molecule counts will be reproducible in this way. Indeed, the example of ion channels makes clear that, in the natural parameter space for the cell, there are many different ways of achieving essentially the same function, and so there is no reason for the cell to control the number of copies of any one particular molecule very precisely; what is important are the tight correlations among variations in different molecule counts, and these correlations are often expected and observed to be nonlinear, even defining non-convex regions of parameter space.

The fact that they can exert precise control over the concentrations, or combinations of concentrations, of certain molecules does not solve all of the organism’s problems. Most fundamentally, life as a cold blooded organism⁷⁵ means having to function across a range of conditions where all chemical reaction rates vary, often by an order of magnitude or more, with no guarantee that the different rates in a given network will scale together; for an example of this problem one need look no further than the familiar circadian rhythms, which have long been known to be invariant to temperature changes. At the same time, diversity of environments is one of the driving forces for speciation, so that (for example) the fruit flies that live at different latitudes, and hence different temperature ranges, are genetically distinguishable. Natural history abounds with stories of animals that seek out very special environments in which to lay their eggs, casting doubt on any glib statement that embryonic development is robust against environmental perturbations. Still, simple laboratory experiments demonstrate that many aspects of life are nearly invariant over a wide range of temperatures, much wider than we might expect from simple models.

Locating life on the spectrum between precisely controlled (rather than finely tuned) dynamics and some more generic or robust behavior is an incredibly important question. It touches, as we have seen, phenomena ranging from the states of single cells to the nature of our memories. It connects to theoretical ideas that have the potential to reach deeply into statistical physics and dynamical systems. Still, at the risk of making clear the limits of my own understanding, I would say that we are still searching for the best formulations of these questions. We need more experimental guidance about what features of behavior are robust against which variations, and we need evidence that organisms actually face these variations in their natural environment. On the theoretical side, we need more anchor points like the random

heteropolymer and the random neural network, where we have a complete analytic understanding of what is expected in the truly generic case, and we need a statistical mechanics of systems with random parameters that allows us to deal with the case where these parameters have nontrivial distributions. These are substantial challenges.

The idea of choosing parameters at random in biochemical networks was explored by Barkai and Leibler (1999) and by von Dassow et al (2000), among others, using simulations; see Sections III.C. Much earlier, Sompolinsky et al (1988) had analyzed the dynamics of random neural networks, identifying a transition between a stationary phase and a chaotic phase at a critical value of the typical synaptic strength. For attempts to connect these random networks to the behavior of cortex, see van Vreeswijk & Sompolinsky (1996, 1998). More recently, Rajan et al (2010) have emphasized that input signals can drive random networks across the transition between order and chaos, providing a possible new view of the nature of variability in cortical responses (Abbott et al 2010).

Abbott et al 2010: Interactions between intrinsic and stimulus-evoked activity in recurrent neural networks. LF Abbott, K Rajan & H Sompolinsky, in *Neuronal Variability and Its Functional Significance*, M Ding & D Glanzman eds, in press (Oxford University Press, Oxford, 2010); arXiv:0912.3832 (2009).

Rajan et al 2009: Input-dependent suppression of chaos in recurrent neural networks. K Rajan, LF Abbott & Sompolinsky, *Phys Rev E* **82**, 011903 (2010).

Sompolinsky et al 1988: Chaos in random neural networks. H Sompolinsky, A Crisanti & HJ Sommers, *Phys Rev Lett* **61**, 259–262 (1988).

van Vreeswijk & Sompolinsky 1996: Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724–1726 (1996).

van Vreeswijk & Sompolinsky 1998: Chaotic balanced state in a model of cortical circuits. *Neural Comp* **10**, 1321–1371 (1998).

There is work on why chemical dynamics tends to be “simple.” Should give pointers here. Maybe also some refs on circadian clocks, and speciation of flies in different latitudes.

⁷⁵ Most of the biomass on our planet is cold blooded, so this is a very general problem.

IV. EFFICIENT REPRESENTATION

The generation of physicists who turned to biological phenomena in the wake of quantum mechanics noted that, to understand life, one has to understand not just the flow of energy (as in inanimate systems) but also the flow of information. There is, of course, some difficulty in translating the colloquial notion of “information” into something mathematically precise. Almost all statistical mechanics textbooks note that the entropy of a gas measures our lack of information about the microscopic state of the molecules, but often this connection is left a bit vague or qualitative. In 1948, Shannon proved a theorem that makes the connection precise: entropy is the unique measure of available information consistent with certain simple and plausible requirements. Further, entropy also answers the practical question of how much space we need to use in writing down a description of the signals or states that we observe. This leads to a notion of *efficient representation*, and in this Chapter we’ll explore the possibility that biological systems in fact form efficient representations, maximizing the amount of relevant information that they transmit and process, subject to fundamental physical constraints.

The idea that a mathematically precise notion of “information” would be useful in thinking about the representation of information in the brain came very quickly after Shannon’s original work. There is, therefore, a well developed set of ideas about the how many bits are carried by the responses of neurons, in what sense the encoding of sensory signals into sequences of action potentials is efficient, and so on. More subtly, there is a body of work on the theory of learning that can be summarized by saying that the goal of learning is to build an efficient representation of what we have seen. In contrast, most discussions of signaling and control at the molecular level has left “information” as a colloquial concept. One of the goals of this Chapter, then, is to bridge this gap. Hopefully, in the physics tradition, it will be clear how the same concepts can be used in thinking about the broadest possible range of phenomena. We begin, however, with the foundations.

A. Entropy and information

Two friends, Max and Allan, are having a conversation. In the course of the conversation, Max asks Allan what he thinks of the headline story in this morning’s newspaper. We have the clear intuitive notion that Max will ‘gain information’ by hearing the answer to his question, and we would like to quantify this intuition. Let us start by assuming that Max knows Allan very well. Allan speaks very proper English, being careful to follow the grammatical rules even in casual conversation. Since they have had many political discussions Max has




FIG. 128 The branching postulate in Shannon’s proof. The idea is to break a big question into multiple parts, as in the familiar game of twenty questions. We start with some initial question, at the top (?). Depending on the answer to this question (*a* or *b*), we ask a new question. This second question in turn has multiple possible answers (*a*₁, *a*₂, *a*₃ or *b*₁, *b*₂). In this tree structure, the various sub-questions live at branch points, with the answers emerging along the branches; finding our way to the full answer means following one path through the tree. The average information that we gain along this path should be additive, the weighted sum of information gained at every branch point.

a rather good idea about how Allan will react to the latest news. Thus Max can make a list of Allan’s possible responses to his question, and he can assign probabilities to each of the answers. From this list of possibilities and probabilities we can compute an entropy, and this is done in exactly the same way as we compute the entropy of a gas in statistical mechanics. Thus, if the probability of the *n*th possible response is *p*_{*n*}, then the entropy is

$$S = - \sum_n p_n \log_2 p_n \text{ bits.} \quad (616)$$

Our intuition from statistical mechanics suggests that the entropy *S* measures Max’s uncertainty about what Allan will say in response to his question, in the same way that the entropy of a gas measures our lack of knowledge about the microstates of all the constituent molecules. Once Allan gives his answer, all of this uncertainty is removed—one of the responses occurred, corresponding to *p* = 1, and all the others did not, corresponding to *p* = 0—so the entropy is reduced to zero. It is appealing to equate this reduction in our uncertainty with the information we gain by hearing Allan’s answer. Shannon proved that this is not just an interesting analogy; it is the *only* definition of information that conforms to some simple constraints.

If we want to have a general measure of how much information is gained on hearing the answer to a question, we have to put aside the details of the questions and the answers—although this might make us uncomfortable,

and is something we should revisit. If we leave out the text of the questions and answers themselves, then all that remains are the probabilities p_n of hearing the different answers, and so Shannon assumes that the information gained must be a function of these probabilities, $I(\{p_n\})$. The challenge is to determine this function.⁷⁶

The first constraint is that, if all N possible answers are equally likely, then the information gained should be a monotonically increasing function of N —we learn more by asking questions that have a wider range of possible answers. The next constraint is that if our question consists of two parts, and if these two parts are entirely independent of one another, then we should be able to write the total information gained as the sum of the information gained in response to each of the two sub-questions. Finally, more general multipart questions can be thought of as branching trees, as in Fig 128, where the answer to each successive part of the question provides some further refinement of the probabilities; in this case we should be able to write the total information gained as the weighted sum of the information gained at each branch point. Shannon proved that the only function of the $\{p_n\}$ consistent with these three postulates—monotonicity, independence, and branching—is the entropy S , up to a multiplicative constant. The proof is sufficiently simple that it seems worth going through the details, not least to be sure we understand how little is required to derive such a powerful result.

To prove Shannon's theorem we start with the case where all N possible answers are equally likely. Then the information must be a function of N , and let this function be $I(\{p_n\}) = f(N)$. Consider the special case $N = k^m$. Then we can think of our answer—one out of N possibilities—as being given in m independent parts, and in each part we must be told one of k equally likely possibilities. But we have assumed that information from independent questions and answers must add, so the function $f(N)$ must obey the condition

$$f(k^m) = mf(k). \quad (617)$$

Notice that although we are focusing on cases where $N = k^m$, we have a condition that involves $f(k)$ for arbitrary k . It is easy to see that $f(N) \propto \log N$ satisfies this equation. To show that this is the unique solution, consider another pair of integers ℓ and n such that

$$k^m \leq \ell^n \leq k^{m+1}, \quad (618)$$

or, taking logarithms,

$$\frac{m}{n} \leq \frac{\log \ell}{\log k} \leq \frac{m}{n} + \frac{1}{n}. \quad (619)$$

Now because the information measure $f(N)$ is monotonically increasing with N , the ordering in Eq. (618) means that

$$f(k^m) \leq f(\ell^n) \leq f(k^{m+1}), \quad (620)$$

and hence from Eq. (617) we obtain

$$mf(k) \leq nf(\ell) \leq (m+1)f(k). \quad (621)$$

Dividing through by $nf(k)$ we have

$$\frac{m}{n} \leq \frac{f(\ell)}{f(k)} \leq \frac{m}{n} + \frac{1}{n}, \quad (622)$$

which is very similar to Eq. (619). The trick is now that with k and ℓ fixed, we can choose an arbitrarily large value for n , so that $1/n = \epsilon$ is as small as we like. Then Eq. (619) is telling us that

$$\left| \frac{m}{n} - \frac{\log \ell}{\log k} \right| < \epsilon, \quad (623)$$

and Eq. (622) for the function $f(N)$ can similarly be rewritten as

$$\left| \frac{m}{n} - \frac{f(\ell)}{f(k)} \right| < \epsilon. \quad (624)$$

Putting these together, we have

$$\left| \frac{f(\ell)}{f(k)} - \frac{\log \ell}{\log k} \right| \leq 2\epsilon, \quad (625)$$

so that $f(N) \propto \log N$ as promised. Note that if we were allowed to consider $f(N)$ as a continuous function, then we could have made a much simpler argument. But, strictly speaking, $f(N)$ is defined only at integer arguments.

We are not quite finished, even with the simple case of N equally likely alternatives, because we still have an arbitrary constant of proportionality. We recall that the same issue arises in statistical mechanics: what are the units of entropy? In a chemistry course you might learn that entropy is measured in “entropy units,” with the property that if you multiply by the absolute temperature (in Kelvin) you obtain an energy in units of calories per mole; this happens because the constant of proportionality is chosen to be the gas constant R , which refers to Avogadro's number of molecules.⁷⁷ In physics

⁷⁶ Notice that Shannon's ‘zeroth’ assumption—that the information gained is a function of the probability distribution over the answers to our question—means that we must take seriously the notion of enumerating the possible answers. In this framework we cannot quantify the information that would be gained upon hearing a literally unimaginable answer to our question. It is interesting to think about whether this is a real restriction.

⁷⁷ I have to admit that whenever I read about entropy units (or calories, for that matter) I imagine that there was some great congress on units at which all such things were supposed to be standardized. Of course every group has its own favorite non-standard units. Perhaps at the end of some long negotiations the chemists were allowed to keep entropy units in exchange for physicists continuing to use electron Volts.

courses entropy is often defined with a factor of Boltzmann's constant k_B , so that if we multiply by the absolute temperature we again obtain an energy (in Joules) but now per molecule (or per degree of freedom), not per mole. In fact many statistical mechanics texts take the sensible view that temperature itself should be measured in energy units—that is, we should always talk about the quantity $k_B T$, not T alone—so that the entropy, which after all measures the number of possible states of the system, is dimensionless. Any dimensionless proportionality constant can be absorbed by choosing the base that we use for taking logarithms, and in measuring information it is conventional to choose base two. Finally, then, we have $f(N) = \log_2 N$. The units of this measure are called *bits*, and one bit is the information contained in the choice between two equally likely alternatives.

Ultimately we need to know the information conveyed in the general case where our N possible answers all have unequal probabilities. Consider first the situation where all the probabilities are rational, that is

$$p_n = \frac{k_n}{\sum_m k_m}, \quad (626)$$

where all the k_n are integers. If we can find the correct information measure for rational $\{p_n\}$ then by continuity we can extrapolate to the general case; the trick is that we can reduce the case of rational probabilities to the case of equal probabilities. To do this, imagine that we have a total of $N_{\text{total}} = \sum_m k_m$ possible answers, but that we have organized these into N groups, each of which contains k_n possibilities, as in Fig 129. To specify the full answer, we would first tell which group it is in, then tell which of the k_n possibilities is realized. In this two step process, at the first step we get the information we




FIG. 129 Grouping. To determine the information gained with unequal probabilities, we consider a “big question” with answer that fall into N groups. By hypothesis, in each the k_n answers are equally likely.

are really looking for—which of the N groups are we in—and so the information in the first step is our unknown function,

$$I_1 = I(\{p_n\}). \quad (627)$$

At the second step, if we are in group n then we will gain $I_n = \log_2 k_n$ bits, because this is just the problem of choosing from k_n equally likely possibilities, and since group n occurs with probability p_n , the *average* information we gain in the second step is

$$I_2 = \sum_n p_n I_n = \sum_n p_n \log_2 k_n. \quad (628)$$

But this two step process is not the only way to compute the information in the enlarged problem, because, by construction, the enlarged problem is just the problem of choosing from N_{total} equally likely possibilities. The two calculations have to give the same answer, so that

$$I_1 + I_2 = \log_2 (N_{\text{total}}), \quad (629)$$

$$I(\{p_n\}) + \sum_n p_n \log_2 k_n = \log_2 \left(\sum_m k_m \right). \quad (630)$$

Rearranging the terms, we find

$$I(\{p_n\}) = - \sum_n p_n \log_2 \left(\frac{k_n}{\sum_m k_m} \right) \quad (631)$$

$$= - \sum_n p_n \log_2 p_n. \quad (632)$$

Again, although this is worked out explicitly for the case where the p_n are rational, it must be the general answer if the information measure is continuous. So we are done: the average information gained on hearing the answer to a question is measured uniquely by the entropy of the distribution of possible answers.

It is worth pausing here to note that what Shannon did is very different from our conventional experience in using mathematics to describe the natural world. In most of physics, we have some set of observations (the motion of the planets in the night sky, for example) that can be made quantitative (as Brahe did), and we search for mathematical structures that can explain and unify these data (Kepler, Newton). In contrast, Shannon considered an everyday phenomenon for which we have a colloquial language, and asked if this language itself could be made mathematically precise, without reference to quantitative data. It is remarkable that this actually worked, and that Shannon's construction has, as we will see, so many consequences.

When we try to quantify the information we gain from hearing the answer to a question, it seems natural to think about a discrete set of possible answers. On the other hand, if we think about gaining information from the acoustic waveform that reaches our ears, then there

is a continuum of possibilities. Naively, we are tempted to write

$$S_{\text{continuum}} = - \int dx P(x) \log_2 P(x), \quad (633)$$

or some multidimensional generalization. The difficulty is that probability distributions for continuous variables have units— $P(x)$ has units inverse to the units of x —and we should be worried about taking logs of objects that have dimensions. Notice that if we wanted to compute a difference in entropy between two distributions, this problem would go away. This is a hint that only entropy differences are going to be important.

Problem 127: Dimensionality and the scaling of the entropy. As written, Eq (633) doesn't really make sense, because we are taking the log of something with units. Suppose we try to clean this up, and make bins along the x axis, each bin of width Δx and the n^{th} bin centered at x_n . Then if the bins are reasonably small, the probability of falling in the n^{th} bin is $p_n = P(x_n)\Delta x$.

(a.) Show that if you calculate the entropy in the usual way, you find

$$S = - \sum_n p_n \log_2 p_n = S_{\text{continuum}} - \log_2(\Delta x) \quad (634)$$

in the limit $\Delta x \rightarrow 0$. More generally, show that in D dimensions

$$S = - \sum_n p_n \log_2 p_n = S_{\text{continuum}} - D \log_2(\Delta x). \quad (635)$$

The result in Eq (635) suggests that the scaling of the entropy with bin size provides a measure of the dimensionality D of the underlying space. This is especially interesting if the intrinsic dimensionality is different from the dimensionality we happen to be using in describing the system. As an example, if we describe a system by its position in a two dimensional space (x, y) , but really the points fall on a curve, then the right answer is that the system is one dimensional, not two dimensional.

(b.) Write a small program in MATLAB to generate 10^6 points in the (x, y) plane that fall on the circle $x^2 + y^2 = 1$. Then divide the plane (you can confine your attention to the region $-2 < x < 2$, and similarly for y) into boxes of size $(\Delta x) \times (\Delta x)$, and estimate the fraction of points that fall in each box. From this estimate, compute the entropy, and see how it varies as a function of Δx . Can you identify the signature of the reduced dimensionality?

(c.) Suppose that you take the 10^6 points from (b) and add, to each point, a bit of noise in the x and y directions, for example Gaussian noise with a standard deviation of $\sigma = 0.05$. Repeat the calculation of the entropy vs. box size. If you look closely enough ($\Delta x \ll \sigma$) the underlying probability distribution really is two dimensional, since there is independent noise along x and y . But if your resolution is more coarse ($\Delta x \gg \sigma$) you won't be able to "see" the noise and the points will appear to fall on a circle, corresponding to a one dimensional distribution. Can you see this transition in the plot of $S(\Delta x)$?

The problem of defining the entropy for continuous

variables is familiar in statistical mechanics.⁷⁸ In the simple example of an ideal gas in a finite box, we know that the quantum version of the problem has a discrete set of states, so that we can compute the entropy of the gas as a sum over these states. In the limit that the box is large, sums can be approximated as integrals, and if the temperature is high we expect that quantum effects are negligible and one might naively suppose that Planck's constant should disappear from the results; we recall that this is not quite the case. Planck's constant has units of momentum times position, and so is an elementary area for each pair of conjugate position and momentum variables in the classical phase space; in the classical limit the entropy becomes (roughly) the logarithm of the occupied volume in phase space, but this volume is measured in units of Planck's constant. If we start with a classical formulation (as did Boltzmann and Gibbs, of course) then we would find ourselves with the problems of Eq. (633), namely that we are trying to take the logarithm of a quantity with dimensions. If we measure phase space volumes in units of Planck's constant, then all is well. The important point is that the problems with defining a purely classical entropy do *not* stop us from calculating entropy differences, which are observable directly as heat flows, and we shall find a similar situation for the information content of continuous ("classical") variables.

In the simple case where we ask a question and there are exactly $N = 2^m$ possible answers, all with equal probability, the entropy is just m bits. But if we make a list of all the possible answers we can label each of them with a distinct m -bit binary number: to specify the answer all we need to do is write down this number. Note that the answers themselves can be very complex—different possible answers could correspond to lengthy essays, but the number of pages required to write these essays is irrelevant. If we agree in advance on the set of possible answers, all we have to do in answering the question is to provide a unique label. If we think of the label as a 'code word' for the answer, then in this simple case the length of the code word that represents the n^{th} possible answer is given by $\ell_n = -\log_2 p_n$, and the average length of a code word is given by the entropy.

The equality of the entropy and the average length of code words is much more general than our simple example. Before proceeding, however, it is important to realize that the entropy is emerging as the answer to two

⁷⁸ Indeed, this problem is so troublesome that it has led to a serious shift in our teaching. It is simpler to define everything in the case where states are discrete, and this has led many people to argue that we shouldn't teach statistical physics until after students have learned quantum mechanics. Whatever advantages this might have, it guarantees that many US students never see anything statistical (beyond a few lectures on the kinetic theory of gases) until their third year of university, which is quite late.

very different questions. In the first case we wanted to quantify our intuitive notion of gaining information by hearing the answer to a question. In the second case, we are interested in the problem of *representing* this answer in the smallest possible space. It is quite remarkable that the only way of quantifying how much we learn by hearing the answer to a question is to measure how much space is required to write down the answer.

Clearly these remarks are interesting only if we can treat more general cases. Let us recall that in statistical mechanics we have the choice of working with a microcanonical ensemble, in which an ensemble of systems is distributed uniformly over states of fixed energy, or with a canonical ensemble, in which an ensemble of systems is distributed across states of different energies according to the Boltzmann distribution. The microcanonical ensemble is like our simple example with all answers hav-

ing equal probability: entropy really is just the log of the number of possible states. On the other hand, we know that in the thermodynamic limit there is not much difference between the two ensembles. This suggests that we can recover a simple notion of representing answers with code words of length $\ell_n = -\log_2 p_n$ provided that we can find a suitable analog of the thermodynamic limit.

Imagine that instead of asking a question once, we ask it many times. As an example, every day we can ask the weatherman for an estimate of the temperature at noon the next day. Now instead of trying to represent the answer to one question we can try to represent the whole stream of answers collected over a long period of time. Let us label the sequences of answers $n_1 n_2 \cdots n_N$, and these sequences have probabilities $P(n_1 n_2 \cdots n_N)$.⁷⁹ From these probabilities we can compute an entropy that must depend on the length of the sequence,

$$S(N) = - \sum_{n_1} \sum_{n_2} \cdots \sum_{n_N} P(n_1 n_2 \cdots n_N) \log_2 P(n_1 n_2 \cdots n_N). \quad (636)$$

Now we can draw on our intuition from statistical mechanics. The entropy is an extensive quantity, which means that as N becomes large the entropy should be proportional to N ; more precisely we should have

$$\lim_{N \rightarrow \infty} \frac{S(N)}{N} = \mathcal{S}, \quad (637)$$

where \mathcal{S} is the entropy density for our sequence in the same way that a large volume of material has a well defined entropy per unit volume.

The equivalence of ensembles in the thermodynamic limit means that having unequal probabilities in the Boltzmann distribution has almost no effect on anything we want to calculate. In particular, for the Boltzmann distribution we know that, state by state, the log of the probability is the energy and that this energy is itself an extensive quantity. Further we know that (relative) fluctuations in energy are small. But if energy is log probability, and relative fluctuations in energy are small, this must mean that almost all the states we actually observe have log probabilities which are the same. By analogy, all the long sequences of answers must fall into two groups: those with $-\log_2 P \approx N\mathcal{S}$, and those with $P \approx 0$. Now this is all a bit sloppy, but it is the right idea: if we are willing to think about long sequences or streams of data, then the equivalence of ensembles tells us that ‘typical’

sequences are uniformly distributed over $\mathcal{N} \approx 2^{N\mathcal{S}}$ possibilities, and that this approximation becomes more and more accurate as the length N of the sequences becomes large.

Problem 128: Probabilities and the equivalence of ensembles.⁸⁰ Consider an ideal monatomic gas in three dimensions, for which the energy is

$$E = \frac{1}{2m} \sum_{i=1}^{3N} p_i^2, \quad (638)$$

where m is the atomic mass. We will define the classical sum over states to be an integral over positions and velocities, normalized by appropriate powers of Planck’s constant h .

(a.) The partition function in the microcanonical ensemble is

$$Z_{\text{micro}}(E) \equiv \frac{1}{h^{3N}} \int d^3x \int d^3p \delta \left(E - \frac{1}{2m} \sum_{i=1}^{3N} p_i^2 \right) \quad (639)$$

$$= \left(\frac{V}{h^3} \right)^N \int d^3p \delta \left(E - \frac{1}{2m} \sum_{i=1}^{3N} p_i^2 \right). \quad (640)$$

If the energy is fixed with precision ϵ , then $Z_{\text{micro}}(E)\epsilon$ is the number of accessible states, all occurring with equal probability, and so the microcanonical entropy is $S_{\text{micro}}(E) = \log_2[Z_{\text{micro}}(E)\epsilon]$. Use the Fourier representation of the delta function and the method of

⁷⁹ Notice that, at this point, we do not need to assume that successive questions have independent answers.

⁸⁰ This should be a review of things you learned in a statistical mechanics class, though perhaps in slightly different language. It is useful to make all of this explicit here.

steepest descent to derive the asymptotic behavior of $S_{\text{micro}}(E)$ at large N .

(b.) In the canonical ensemble, at inverse temperature β , the probability of being in any state is given by the Boltzmann distribution,

$$P = \frac{1}{Z(\beta)} e^{-\beta E}, \quad (641)$$

where

$$Z(\beta) = \frac{1}{h^{3N}} \int d^3x \int d^3p \exp\left(-\frac{\beta}{2m} \sum_{i=1}^{3N} p_i^2\right). \quad (642)$$

Evaluate $Z(\beta)$ and the entropy $S(\beta)$. Review what we mean when we say that the entropy is the same in the canonical and micro-canonical ensembles at large N .

(c.) The typical probability of a state in the canonical ensemble is $P_{\text{typical}} = 2^{-S(\beta)}$. Define the deviation from this typical probability as $\Delta = \log_2(P/P_{\text{typical}})$. What can you say about the distribution of Δ over all the states? Can you make a precise version of the statement that “most” states have either “almost” the typical probability or zero probability? For example, can you put a bound on the fraction f of states which have $|\Delta| > \delta_c$? How does the relation between f and δ_c change with N ?

Problem 129: More about typicality. Consider drawing N samples of a variable that can take on K different values, with probabilities p_1, p_2, \dots, p_K . Let the sequence of samples that you observe be called i_1, i_2, \dots, i_N , which has probability

$$P = \prod_{n=1}^N p_{i_n}. \quad (643)$$

It should be easy to show that the average of $L = -(1/N) \log_2 P$ is the entropy of the underlying distribution, $S = -\sum_i p_i \log_2 p_i$. Say as much as you can about the distribution of L as N becomes large.

The idea of typical sequences, which is the information theoretic version of a thermodynamic limit, is enough to tell us that our simple arguments about representing answers by binary numbers ought to work on average for long sequences of answers. An important if obvious consequence is that if we have many rather unlikely answers (rather than fewer more likely answers) then we need more space to write the answers down. More profoundly, this turns out to be true answer by answer: to be sure that long sequences of answers take up as little space as possible, we need to use $\ell_n \approx -\log_2 p_n$ bits to represent each individual answer n . Thus, even individual answers which are more surprising require more space to write down.

As a simple example, imagine that we have four answers, with probabilities $p_1 = 1/2$, $p_2 = 1/4$, and $p_3 = p_4 = 1/8$. Naively, if we use a binary representation we will need two bits to represent the four possibilities. But the entropy is

$$S \equiv \sum_{i=1}^4 p_i \log_2 p_i = \frac{1}{2} \log_2 2 + \frac{1}{4} \log_2 4 + \frac{2}{8} \log_2 8 = \frac{7}{4}, \quad (644)$$

which is less than two bits (as it must be). Suppose that we represent the four possibilities by the binary sequences:

$$1 \rightarrow 0, \quad (645)$$

$$2 \rightarrow 10, \quad (646)$$

$$3 \rightarrow 110, \quad (647)$$

$$4 \rightarrow 111. \quad (648)$$

Notice that the length of each code word obeys $\ell_i = -\log_2 p_i$, so we know that, on average, the number of binary digits that we use per answer will be equal to the entropy. This illustrates the idea that, by using code words of different lengths, we can reduce the average amount of space we need to write things down.

Problem 130: Do we need commas? When we represent a sequence of answers, we have to be sure that we can find the boundaries between the code words. If all the words have the same length, we can just count, but this doesn't work if we use unequal lengths. At worst, we could add an extra symbol to “punctuate” the stream of words, but this takes extra space and surely is inefficient. Convince yourself that the code defined by Eqs (645) through (648) does not need any extra symbols—all sequences of code words can be parsed uniquely.

To complete the picture, we have to put together the ideas of typicality and code words of varying length. Suppose that we look at a block of N answers, n_1, n_2, \dots, n_N as before; let's label this block (or “state,” to reinforce the analogy with statistical physics) by s , which occurs with probability p_s . We choose the labels so that all the states are numbered in order of their probability, that is $p_1 \geq p_2 \geq \dots \geq p_K$, where K is the number of possible sequences of length N . For each state s we can compute the cumulative probability of lower energy (higher probability) states, $P_s \equiv \sum_{i=1}^{s-1} p_i$. Now take this cumulative probability and expand it as a binary number. If we stop after m_s digits, where

$$-\log_2 p_s \leq m_s < -\log_2 p_s + 1, \quad (649)$$

then we guarantee that this binary number we are looking at will be different from any subsequent number with larger s , so it is a unique encoding of the state s , as shown schematically in Fig 130. But now we can see that the average number of binary digits we have used to encode the blocks of length N will be

$$L(N) \equiv \sum_s p_s m_s, \quad (650)$$




FIG. 130 Coding of sequences with variable word length. In a stream where ‘0’ and ‘1’ occur independently, but with unequal probabilities, we can compress our description by coding N -bit blocks; here $N = 10$. Each block can be labelled by s , the number equivalent to the binary string (top). These states have widely varying probability p_s (lower left). We can compute the cumulative probability of states with higher probability (lower right), as described in the text, and use the binary expansion of this cumulative probability as the code word W . We stop the expansion at a number of digits given by rounding up from $-\log_2(p_s)$.

and we can bound this from both sides,

$$\sum_s p_s m_s (-\log_2 p_s) \leq L(N) < \sum_s p_s (-\log_2 p_s + 1) \tag{651}$$

$$S(N) \leq L(N) < S(N) + 1, \tag{652}$$

where S is the entropy of the N -answer blocks, $S(N) = -\sum_s p_s \log_2 p_s$. If we count the length of the code *per answer*, then

$$\frac{S(N)}{N} \leq \frac{L(N)}{N} < \frac{S(N)}{N} + \frac{1}{N}. \tag{653}$$

But, as before, we know that the entropy per degree of freedom should approach a finite entropy density, as in Eq (637), and now we see that the average code length per answer is within $1/N$ of this entropy density. Thus, as $N \rightarrow \infty$, the entropy and the minimum code length are equal.

To summarize, if we need to write down answers many times, then the minimum space required to write down these answers is, per answer, the entropy of the distribution out of which the answers are drawn. Notice that our choice of alphabet in which to write is arbitrary, but we also had an arbitrariness in choosing the units of entropy; this is the same arbitrariness. Thus, the statement that entropy is both the amount of information we gain

and the amount of space we need to write down what we have learned is *not* arbitrary, and there are no constants floating around to spoil the exact equality. To reach this maximally compact representation, we must at least implicitly use the structure of the probability distribution out of which the answers are drawn, adjusting the lengths of individual code words in relation to the probability of the answer.

Problem 131: Coding rare events. Suppose that we have two possible answers, A and B , which occur with very unequal probabilities, $p_A \ll p_B$. Show that the entropy of the distribution of answers is approximately $S \approx p_A \log_2(e/p_A)$. If we have a long sequence of answers, most are B with a sprinkling of A s. Try to encode such a sequence in binary form, using a code in which some symbol (e.g., 1111) is reserved for A , and the blocks of B are encoded by writing the number of consecutive B s as a binary number. To make this work—that is, to be sure that your encoding can be uniquely decoded—you obviously have to be careful in the special case where the number of B s is equal to 15 (1111 in binary form). Are there any other problems? Can you find a solution? Does this code come close to the lower bound on code length set by the entropy?

The idea that there is a minimum amount of space required to write down a description of a system is incredibly important. At a practical level, we pay for the resources needed to write things down, or to transmit information from one place to another, and so there is a premium on using as little space as possible. This is often called “data compression.” More generally, this is the first indication that there is a general notion of efficiency in representing data, and we will see how this becomes relevant to biological systems.

The argument we have just given tells us that once we know the probability distribution for the states s , we have a code that we can use to represent these states, and asymptotically this code is of minimum length. Suppose that states really are chosen out of a distribution $\mathbf{p} \equiv \{p_s\}$, but we don’t know this; instead, we think that the distribution is \mathbf{q} . Then (neglecting terms that are unimportant in the large N limit), we assign a code word of length $\ell_s = -\log_2 q_s$ to each state, and so the mean code length is

$$L = -\sum_s p_s \log_2 q_s. \tag{654}$$

This is different than the entropy of the distribution \mathbf{p} , and the difference

$$L - L_{\min} = L - S = - \sum_s p_s \log_2 q_s - \left[- \sum_s p_s \log_2 p_s \right] = \sum_s p_s \log_2 \left(\frac{p_s}{q_s} \right). \quad (655)$$

This quantity is zero if the two distributions are the same, and is positive for any pair of distributions \mathbf{p} and \mathbf{q} ; it is called the Kullback–Leibler divergence between the two distributions, and usually is written as

$$D_{KL}(\mathbf{p}||\mathbf{q}) = \sum_s p_s \log_2 \left(\frac{p_s}{q_s} \right). \quad (656)$$

Notice that this is not a symmetric quantity, and hence is not a metric on the space of distributions, although it does say something about the degree of similarity or difference between \mathbf{p} and \mathbf{q} . D_{KL} also is sometimes called the “relative entropy” of the distribution \mathbf{p} with respect to \mathbf{q} .

To emphasize the role of D_{KL} as a measure of difference between distributions, suppose that we are given N samples and have to decide whether they came from \mathbf{p} or \mathbf{q} . Out of the N samples, n_1 come from state 1, n_2 come from state 2, and so on. So the probability that the distribution \mathbf{p} generated these samples is

$$P(\text{samples}|\mathbf{p}) = A \prod_s p_s^{n_s}, \quad (657)$$

where A is a combinatorial factor, and similarly

$$P(\text{samples}|\mathbf{q}) = A \prod_s q_s^{n_s}. \quad (658)$$

What we want to know is, given the samples, what is the probability P that they came from the distribution \mathbf{p} as opposed to \mathbf{q} ? Let us say that, a priori, the two possibilities are equally likely. Then, by Bayes’ rule,

$$P = \frac{P(\text{samples}|\mathbf{p})P(\mathbf{p})}{P(\text{samples})} \quad (659)$$

$$= \frac{P(\text{samples}|\mathbf{p})}{P(\text{samples}|\mathbf{p}) + P(\text{samples}|\mathbf{q})} \quad (660)$$

$$= \frac{1}{1 + 2^{-\Lambda}}, \quad (661)$$

where

$$\Lambda = \log_2 \left[\frac{P(\text{samples}|\mathbf{p})}{P(\text{samples}|\mathbf{q})} \right] = \sum_s n_s \log_2 \left(\frac{p_s}{q_s} \right). \quad (662)$$

As discussed in Chapter 1 [give specific pointer], Λ is called the log likelihood ratio. We notice that since it is proportional to all the n_s , it must also be proportional to N , and hence grows (on average) linearly with the number of samples. We can think of this as the accumulation of evidence for \mathbf{p} vs. \mathbf{q} , and the rate at which this

evidence accumulates is, asymptotically,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \Lambda = \sum_s \left[\lim_{N \rightarrow \infty} \frac{n_s}{N} \right] \log_2 \left(\frac{p_s}{q_s} \right) \quad (663)$$

$$= \sum_s p_s \log_2 \left(\frac{p_s}{q_s} \right) \quad (664)$$

$$= D_{KL}(\mathbf{p}||\mathbf{q}). \quad (665)$$

Thus, the Kullback–Leibler divergence is, like the entropy itself, the answer to two very different questions: the cost of coding data using codes based on the wrong distribution, and the ease of discriminating the distributions from one another based on samples.

Problem 132: A little more about the Kullback–Leibler divergence.

(a.) Show that $D_{KL}(\mathbf{p}||\mathbf{q})$ is positive (semi-)definite, and is minimized when $\mathbf{p} = \mathbf{q}$.

(b.) $D_{KL}(\mathbf{p}||\mathbf{q})$ is unbounded, so some probability distributions are infinitely different from one another. Explain, using the connection to the accumulation of evidence, how to make sense out of this divergence.

(c.) If we have a family of distributions that depend on a parameter, \mathbf{p}_θ , show that $D_{KL}(\mathbf{p}_\theta||\mathbf{p}_{\theta'})$ behaves as $F(\theta) \times (\theta - \theta')^2$ when the parameters θ and θ' are close. Give an explicit formula for $F(\theta)$.

(d.) Imagine that we draw N samples out of the distribution p_{θ_0} , but all we know is that the distribution is in the family p_θ . Use Bayes’ rule to construct $P(\theta|\text{samples})$, and show that as N becomes large this becomes peaked around the right answer, $\theta = \theta_0$. Show that the variance around this peak is related to $F(\theta_0)$.

(e.) If the two distributions \mathbf{p} and \mathbf{q} are Gaussians, it’s relatively easy to evaluate $D_{KL}(\mathbf{p}||\mathbf{q})$. Suppose that the two Gaussians differ in either their means or their variances, but not both. You should find that the choice of changing mean vs. variance makes a difference to the (a)symmetry of D_{KL} . Make this explicit, and use what we have shown about D_{KL} as a measure of discrimination to explain the origin of this difference.

The connection between entropy and information has (at least) one more very important consequence: correlations or order reduce the capacity to transmit information. Perhaps the most familiar example is in spelling. If all possible combinations of letters were legal words, then there would be $(26)^4 = 456,976$ four letter words. But if you look through a large, reasonably coherent body of English text—the collected works of a prolific author, or the last year of newspaper articles—you will find that there at most a few hundred four letter words being used. Most of this restriction of vocabulary comes from correlations among the letters in the word: once we have put

a ‘t’ in the first position, it is much more likely that we will put a vowel in the second position; if we want to put a consonant then it has a high probability of being an ‘h’, and so on. It is important that, while correlations have signs—we speak both of correlation and anti-correlation—with respect to the entropy all correlations have the same effect, namely reducing the entropy. Indeed, as explained in [pointer to appendix on maximum entropy], we can construct models for the probability distribution of the states in a system that are consistent with some measured correlations but otherwise have the maximum possible entropy, and we can build a hierarchy of these models with ever smaller entropies as we take account of more correlations; once we capture all the relevant correlations, the entropy converges to its true value.

For four letter words, as an example, the entropy for random letters would be $S_{\text{rand}} = 4 \log_2(26) = 18.8$ bits. In the collected works of Jane Austen, the “one body” correlations, which measure unequal frequencies with which letters are used, reduces this to $S_{\text{ind}} = 14$ bits. Taking account of the “two body” correlations between pairs of letters cuts this entropy nearly in half, to $S_2 = 7.48$ bits, while the true entropy of the distribution of four letter words in these texts is only slightly less, at $S = 6.92$ bits. Thus the entropy is nearly reduced by a factor of three from the case of completely random letters, and most of this reduction is explained by one and two body correlations. Again, the important point is that these correlations, which may have many advantages, certainly have the consequence of reducing our vocabulary and hence our capacity to transmit information.

This seems an appropriate moment to recall that entropy is a very old idea. It arises in thermodynamics first as a way of keeping track of heat flows, so that a small amount of heat dQ transferred at absolute temperature T generates a change in entropy $dS = dQ/T$. While there is no function Q which measure the heat content of a system, there is a function S that characterizes the (macroscopic) state of a system independent of the path to that state. But now we know that the entropy of a probability distribution also measures the amount of space that we need to write down a description of the (microscopic) states drawn out of that distribution.

[Would a schematic help here?] Let us imagine, then, a thought experiment in which we measure (with some fixed resolution) the positions and velocities of all the gas molecules in a small box, and type these numbers into a file on our computer. There are relatively efficient programs (gzip, or “compress” on a UNIX machine) that compresses such files to nearly their shortest possible length. If this really works, then the length of the file tells us the entropy of the distribution out of which the numbers in the file are being drawn, but this is the entropy of the gas. Thus, if we heat up the room by ten degrees, and repeat the process, we will find that the

resulting data file is longer. More profoundly, if we measure the increase in the length of the file, we know the entropy change of the gas and hence the amount of heat that we had to add to the room in order to increase the temperature. This connection between a rather abstract quantity such as the length, in bits, of a computer file and a very tangible physical quantity such as the amount of heat added to a room has long struck me as one of the more dramatic, if elementary, examples of the power of mathematics to unify our description of very disparate phenomena.

Problem 133: Heat flows and file sizes. Give a problem that expands the thought experiment in the previous paragraph ... maybe with a polymer and entropic forces, where we can simulate?

Returning to the conversation between Max and Allan, we assumed that Max would receive a complete answer to his question, and hence that all his uncertainty would be removed. This is an idealization, of course. The more natural description is that, for example, the world can take on many states w , and by observing data d we learn something but not everything about w . Before we make our observations, we know only that states of the world are chosen from some distribution $P(w)$, and this distribution has an entropy $S[P(w)]$. Once we observe some particular datum d , our (hopefully improved) knowledge of w is described by the conditional distribution $P(w|d)$, and this has an entropy $S[P(w|d)]$ that is smaller than $S[P(w)]$ if we have reduced our uncertainty about the state of the world by virtue of our observations. We identify this reduction in entropy as the information that we have gained about w ,

$$I(d \rightarrow w) \equiv S[P(w)] - S[P(w|d)]. \quad (666)$$

Notice that this depends on exactly what datum d we have observed.

Before proceeding, I should draw attention to some notational issues. Strictly speaking, entropy is a property of the probability distribution out of which the states of a system are drawn. Thus, we write $S[P(w)]$ to mean the entropy of the states of the world when these are drawn out of $P(w)$. Similarly, we should write $S[P(w|d)]$ for the entropy of states of the world conditional on having observed the data d . Notice that $S[\dots]$ is the same functional in both cases. But, this is slightly cumbersome. Indeed, in statistical mechanics and thermodynamics we seldom talk about “the entropy of the distribution out of which the states of the gas have been drawn” (although we should); instead we just say “the entropy of the gas.” In this spirit, sometimes I will write in the shorthand

$S(w) \equiv S[P(w)]$, and $S(w|d) \equiv S[P(w|d)]$. I hope this doesn't cause any confusion.

There is one more notational difficulty. When we talk about the states w of the world, it is natural to say that these states are drawn from the distribution $P(w)$. Similarly, when we talk about the data that we will collect, it is natural to write that particular observations d are drawn from the distribution $P(d)$. The problem is that $P(\cdot)$ refers to different functions in these two cases. We could solve this by noting carefully that the states of the world w come from a set of possible states, $w \in W$, and the distribution over these states should be written $P_W(w)$. Similarly, individual observations come from a set of possible observations, $d \in D$, and the distribution of these data should be written $P_D(d)$. Whenever there is a possibility for confusion, I'll try to adhere to this convention. In other cases, I'll slide to the more informal $P(w)$ and $P(d)$. Again, I hope this doesn't cause problems. **[I am not sure that the current draft lives up to this policy, so please read carefully!]**

With the notational issues settled, let's go back to our problem. Having defined the information gained in Eq (??), we should appreciate that this is not guaranteed

to be positive. Consider, for instance, data which tell us that all of our previous measurements have larger error bars than we thought: clearly such data, at an intuitive level, reduce our knowledge about the world and should be associated with a negative information. Another way to say this is that some data points d will increase our uncertainty about state w of the world, and hence for these particular data the conditional distribution $P(w|d)$ has a larger entropy than the prior distribution $P(w)$, so that I_d will be negative. On the other hand, we hope that, on average, gathering data corresponds to gaining information: although single data points can increase our uncertainty, the average over all data points does not.

If we average over all possible data—weighted, of course, by their probability of occurrence $P_D(d)$ —we obtain the average information that d provides about w :

$$\langle (d \rightarrow w) \rangle = S(w) - \sum_d P_D(d) S(w|D). \quad (667)$$

This can be rearranged and simplified, and the result is so important that it is worth being very explicit about the algebra:

$$\langle (d \rightarrow w) \rangle = - \sum_w P_W(w) \log_2 P_W(w) - \sum_d P_D(d) \left[- \sum_w P(w|d) \log_2 P(w|d) \right] \quad (668)$$

$$= - \sum_w \sum_d P(w, D) \log_2 P_W(w) + \sum_w \sum_d P(w|D) P_D(d) \log_2 P(w|d) \quad (669)$$

$$= - \sum_w \sum_d P(w, D) \log_2 P_W(w) + \sum_w \sum_d P(w, D) \log_2 P(w|d) \quad (670)$$

$$= \sum_w \sum_d P(w, D) \log_2 \left[\frac{P(w|d)}{P_W(w)} \right] \quad (671)$$

$$= \sum_w \sum_d P(w, D) \log_2 \left[\frac{P(w, d)}{P_W(w) P_D(d)} \right], \quad (672)$$

where we identify the joint distribution of states of the world and data, $P(w, d) = P(w|d) P_D(d)$.

We see that, after all the dust settles, the average information which d provides about w is symmetric in d and w . This means that we can also view the state of the world as providing information about the data we will observe, and this information is, on average, the same as the data will provide about the state of the world. This 'information provided' is therefore often called the mutual information, and this symmetry will be very important in subsequent discussions; to remind ourselves of this symmetry we write $I(d; w)$ rather than $\langle (d \rightarrow w) \rangle$.

One consequence of the symmetry or mutuality of information is that we can write the mutual information as

a difference of entropies if two different ways,

$$I(d; w) = S(w) - \sum_d P_D(d) S(w|d) \quad (673)$$

$$= S(d) - \sum_w P_W(w) S(d|w). \quad (674)$$

If we consider only discrete sets of possibilities then entropies are positive (or zero), so that these equations imply

$$I(d; w) \leq S(w) \quad (675)$$

$$I(d; w) \leq S(d). \quad (676)$$

The first equation tells us that by observing d we cannot learn more about the world than there is entropy in

the world itself. This makes sense: entropy measures the number of possible states that the world can be in, and we cannot learn more than we would learn by reducing this set of possibilities down to one unique state. Although sensible (and, of course, true), this is not a terribly powerful statement: seldom are we in the position that our ability to gain knowledge is limited by the lack of possibilities in the world around us. On the other hand, there is a tradition of studying the biological systems as they responds to highly simplified signals, and under these conditions the lack of possibilities in the world can be a significant limitation, substantially confounding the interpretation of experiments.

Equation (676), however, is much more powerful. It says that, whatever may be happening in the world, we can never learn more than the entropy of the distribution that characterizes our data. Thus, if we ask how much we can learn about the world by taking readings from a wind detector on top of the roof, we can place a bound on the amount we learn just by taking a very long stream of data, using these data to estimate the distribution $P_D(d)$, and then computing the entropy of this distribution.

The entropy of our observations thus limits how much we can learn no matter what question we were hoping to answer, and so we can think of the entropy as setting (in a slight abuse of terminology) the capacity of the data d to provide or to convey information. As an example, the entropy of neural responses sets a limit to how much information a neuron can provide about the world, and we can estimate this limit even if we don't yet understand what it is that the neuron is telling us (or the rest of the brain).

Problem 134: Maximally informative experiments.

Imagine that we are trying to gain information about the correct theory T describing some set of phenomena. At some point, our relative confidence in one particular theory is very high; that is, $P(T = T_*) > F \cdot P(T \neq T_*)$ for some large F . On the other hand, there are many possible theories, so our absolute confidence in the theory T_* might nonetheless be quite low, $P(T = T_*) \ll 1$. Suppose we follow the “scientific method” and design an experiment that has a yes or no answer, and this answer is perfectly correlated with the correctness of theory T_* , but uncorrelated with the correctness of any other possible theory—our experiment is designed specifically to test or falsify the currently most likely theory. What can you say about how much information you expect to gain from such a measurement? Suppose instead that you are completely irrational and design an experiment that is irrelevant to testing T_* but has the potential to eliminate many (perhaps half) of the alternatives. Which experiment is expected to be more informative? Although this is a gross cartoon of the scientific process, it is not such a terrible model of a game like “twenty questions.” It is interesting to ask whether people play such question games following strategies that might seem irrational but nonetheless serve to maximize information gain. Related but distinct criteria for optimal experimental design have been developed in the statistical literature.

[I wonder if I should go through the basic calculation of maximum entropy counting here ... since the “things” we count have a cost, this would complete the thought about bounds. At least need a pointer to Appendix A.8.]




FIG. 131 A schematic of how a train of action potential is converted to discrete “words” at different times resolutions $\Delta\tau$. There is a minimum inter-spike interval, the “refractory period” (here, $\sim 2 \text{ ms}$), so that for sufficiently small $\Delta\tau$ the words are binary. Highlighted is the case where $\Delta\tau = 2 \text{ ms}$ and $T = 8 \text{ ms}$, so this segment of the spike train becomes three successive four bit words, 0101, 0100, and 0110.

To see how the ideas of entropy reduction and information work in a real example, let's consider the response of a neuron to sensory inputs. As we have discussed [start in Chapter One; give specific pointers], most neurons in the brain generate a sequence of brief ($\sim 1 \text{ ms}$), identical electrical pulses called action potentials or spikes. Since these events are identical, we can think of them as marking points in time, and then we can build a discrete vocabulary of responses by fixing some limited time resolution $\Delta\tau$, as in Fig 131. More precisely, if $\Delta\tau$ is small, then in each small time window of duration $\Delta\tau$ we will see either one or zero spikes, and so the response is naturally discrete and binary. Then segments of the spike train of duration T can be thought of as $T/\Delta\tau$ -letter binary words. Recording from a single neuron as the animal experiences some reasonably complex, dynamic sensory inputs, it is relatively easy to estimate the distribution of these these words, $P(W)$, so long as we don't make the ratio $T/\Delta\tau$ too large. Then we can compute the entropy of this distribution, $S(T, \Delta\tau)$.

Figure 132 shows the results of experiments on the motion sensitive neuron H1 in the fly visual system that we met earlier, in Section [**], when we discussed noise and the precision of visual motion estimation. In these experiments, the fly sees a randomly moving pattern, and H1 responds with a stream of spikes. If we fix $\Delta\tau = 3 \text{ ms}$ and look at $T = 30 \text{ ms}$ segments of the spike train, there are $2^{T/\Delta\tau} \sim 10^3$ possible words, but the distribution is




FIG. 132 [Make a new version of this.] A neuron responds to dynamic stimuli with sequences of spikes. In this case, as described in the text, we look at the motion sensitive neuron H1 in the fly’s visual system. (A) Each line across time is a single presentation of a movie, and dots mark the arrival times of spikes on each trial. (B) and (C) show the discretization of the spike trains into binary “words” with $\Delta\tau = 3$ ms resolution, and the distribution of words that occur at a particular moment in the movie, $P(w|t)$. (D) The distribution of words averaged over all times, in rank order. From de Ruyter van Steveninck et al (1997).

strongly biased and the entropy is only $S(T, \Delta\tau) \sim 5$ bits. This relatively low entropy means that we can still sample the distributions of words even out to $T \sim 50 - 60$ ms, which is interesting because the fly can actually generate a flight correction in response to visual motion inputs within ~ 30 ms.

The entropy $S(T, \Delta\tau)$ should be an extensive quantity, which means that, for large T , we should have $S(T, \Delta\tau) \propto T$. More strongly, if the correlations in the spike train are sufficiently short ranged, then we expect that at large T we will have

$$\frac{1}{T}S(T, \Delta\tau) = S(\Delta\tau) + \frac{C(\Delta\tau)}{T} + \dots, \quad (677)$$

where \dots vanish more rapidly than $1/T$. In fact we see this in the real data (Fig 133), which suggests that we really can estimate the entropy rate $S(\Delta\tau)$.

Connecting to the discussion above, the entropy rate $S(\Delta\tau)$ sets a limit on the rate at which the spikes can

provide information about the sensory input. When we make $\Delta\tau$ smaller, the entropy rate necessarily goes up, because previously indistinguishable responses map to different words at higher time resolution. Concretely, if we make $\Delta\tau$ smaller by a factor of two, then every ‘1’ in the coarse words can become either a ‘01’ or a ‘10’ in the higher resolution words, and so we expect the entropy to increase by roughly one bit for every spike, as in Fig 131.

Problem 135: Entropy and entropy rate in simple models. Going back to Chapter 1, you know how to generate events drawn from a Poisson process with an arbitrary time dependent rate $r(t)$. Here you should take this (semi-)seriously as a model for spike trains, and use the resulting simulations to explore the entropy and entropy rate of neural responses.

(a.) Start with $r = r_0$, a constant. Generate a long sequence of spikes (e.g., $\sim 10^4$). Choose a time resolution $\Delta\tau$ such that $r_0\Delta\tau \ll 1$, and turn your simulated spike train into a binary sequence; for simplicity ignore the (rare) occurrence of two spikes




FIG. 133 Entropy is extensive. From the experiments on the neuron H1 in Fig 132, we compute the entropy of words at fixed time resolution $\Delta\tau = 3$ ms and variable length T , stopping when T is so large that we can no longer reliably sample the distribution $P(W)$. The data (error bars are smaller than the symbols) fall on the line predicted in Eq (677), and we can thus extract an estimate of the entropy rate $\mathcal{S}(\Delta\tau)$. Redrawn from Strong et al (1998a).

in one bin. Form “words” with $T/\Delta\tau$ bits, and estimate the distribution of these words from your simulated data. Compute the entropy of this distribution, and explore its dependence on T , r_0 , and $\Delta\tau$. Do you see the emergence of an entropy rate, $S \sim ST$?

(b.) Explain why, for a Poisson process with a constant rate, $S = \mathcal{S}T$ should be exact. From this result, you can calculate \mathcal{S} by thinking about just one bin of size $\Delta\tau$, and you should do this. How does your analytic result compare with the simulation results in (a)?

(c.) Suppose that $x(t)$ is a Gaussian stochastic process with correlation function $\langle x(t)x(t') \rangle = \sigma^2 e^{-|t-t'|/\tau_c}$. [This should be explained somewhere already!] Samples of this process can be generated by simulating the Langevin equation,

$$\tau_c \frac{dx}{dt} = -x + 2\sigma\eta(t), \quad (678)$$

where $\langle \eta(t)\eta(t') \rangle = \delta(t-t')$. Consider a Poisson process with rate $r(t) = r_0 e^{x(t)}$. Generate spike sequences for this process, and follow the procedures in (a) to estimate the entropy in binary words of duration T at resolution $\Delta\tau$, with reasonable choices of parameters. Can you observe the emergence of extensive behavior, $S \sim ST$? Does this (as seems plausible) require $T \gg \tau_c$? How do your results depend on σ ?

A long standing question in thinking about the brain has been whether the precise timing of individual spikes is important, or whether the brain is capable of counting spikes only in relatively coarse time bins, so that the “rate” of spikes over longer periods of time is all that matters. We now have the tools to give a more precise formulation of this question. As we increase our time resolution, the entropy of the spike trains goes up, and

hence so does the capacity of the neuron to convey information. The question is whether this capacity is used—does the information about sensory inputs also rise as the time resolution is improved, or is the extra entropy just ‘noise’?

If the sensory inputs are called s , then the information that the spike sequences in some window T provide about these inputs can be written, as in Eq (674), as a difference of entropies,

$$I(s; W) = S(W) - \langle S(W|s) \rangle_s, \quad (679)$$

where $\langle \dots \rangle_s$ denotes an average over the distribution of inputs. We have already discussed the entropy of the neural vocabulary, $S(W)$; the problem is how to estimate $S(W|s)$, the entropy of the words given the sensory input s . To do this we need to sample the distribution $P(W|s)$, that is the distribution of neural responses when the stimulus is fixed. At a minimum, this requires that we repeat the same stimuli many times. So, if the visual stimulus is a long movie, we have to show the movie over and over again. But how do we pick out a particular stimulus s from the continuous stream? One way to do this is to realize that the flow of time in the movie provides an index into the stimuli, and all we need is to be able to compute averages over the distribution of stimuli. If the source of stimuli is ergodic (which we can arrange to be true in the lab!), then an average over stimuli is equivalent to an average over time. So, if we repeat the movie many times, and focus on events at time t relative to the start of the movie, we can sample, in repeats of the movie, the distribution $P(W|t)$, as in Fig 132, and hence estimate $S(W|t)$. Finally, the information is obtained by explicitly replacing the ensemble average with a time average,

$$I(s; W) = S(W) - \langle S(W|t) \rangle_t. \quad (680)$$

Each of the entropy terms on the right should behave as in Eq (677), and so we can extract an estimate of the information rate $R_{\text{info}}(\Delta\tau)$ as a function of time resolution. Results are shown in Fig 134.

We see that, as we vary the time resolution from 800 ms down to 2 ms, the information rate follows the entropy rate, with a nearly constant 50% efficiency. Although we should not generalize too much from one example, this certainly suggests that neurons are making use of a significant fraction of their capacity in actually encoding sensory signals. Also, this is true even at millisecond time resolution. The idea that the entropy of the spike train sets a limit to neural information transmission emerged almost immediately after Shannon’s work, but it was never clear whether these limits could be approached by real systems.




FIG. 134 Entropy and information in a spike train. Experiments on the fly’s motion sensitive visual neuron were analyzed as described in the text (following Fig 132) to estimate the total entropy and the information carried about the sensory input. As we vary the time resolution of our analysis from $\Delta\tau = 800$ ms down to $\Delta\tau = 2$ ms, we distinguish finer details of the neural response and expand the capacity of the putative neural code; this enhanced capacity is measured by the increasing entropy. Remarkably, across this huge range, capacity is used with almost constant efficiency. From Strong et al (1998a).

Problem 136: Information from single events. This section began by defining the information gained in a single observation. Here, we would like to give the parallel for individual neural responses, but there is a twist because spikes are rare compared with silences. Thus it makes sense to ask how much information we obtain per spike, or per non-silent word W . Imagine that we look in a window of duration $\Delta\tau$ at time t , and we are looking for some event e —this event could be a single action potential, or some combination of multiple spikes with specific intervals between them. On average these events occur with some rate \bar{r}_e .

(a.) In the small window $\Delta\tau$, either the event e occurs or it does not; for sufficiently small $\Delta\tau$, the probability of occurrence is $p_e = \bar{r}_e\Delta\tau$. What is the entropy of the binary variable marking the occurrence of the event? Can you simplify your result when $p_e \ll 1$? You’ll see that the entropy in this limit is small, but so is the expected number of events. What is the entropy per event?

(b.) If we know the sensory inputs to this neuron, then the probability of an event depends on time, locked to the time dependence of the sensory signal. Let’s call the time dependent rate $r_e(t)$. As in (a.), compute the entropy of the binary event/nonevent variable, but now conditional on knowledge of the sensory inputs.

(c.) Combine your results in (a.) and (b.) to give an expression for the mean information that the occurrence or non-occurrence of the event provides about the sensory input. Normalize by the expected number of events, to give bits per event. Is the limit $\Delta\tau \rightarrow 0$ well behaved? When the dust settles, you should find that the information per event is

$$I_e = \left\langle \frac{r_e(t)}{\bar{r}_e} \log_2 \left[\frac{r_e(t)}{\bar{r}_e} \right] \right\rangle_t. \quad (681)$$

(d.) As an alternative view of the same question, suppose that we observe a large window of time T . If T is sufficiently large, we can be sure that the event e will occur, but we don’t know when.

Problem 137: Information from single spikes in a simple model. In Problem [**] above, you constructed a model spike train using a Poisson process with a time varying rate $r(t) = r_0 e^{x(t)}$, where $x(t)$ is a Gaussian stochastic process. Show that, for this

model, the information carried by a single spike about $x(t)$ is linear in the variance of the signal $\langle x^2 \rangle$. This suggests that if the signal variance grows, the information carried by spikes grows with it, without bound. Explain what is wrong with this picture. Suppose instead that the spike rate $r(t)$ depends on x through some saturating function, for example

$$r(t) = \frac{r_0}{1 + \exp[-x(t) + \theta]}. \quad (682)$$

Reduce the formula for I_e in this model to a single integral which you can do numerically. Can you see how the results simplify as $\langle x^2 \rangle$ becomes large? As a hint, notice that this is equivalent to a model in which

$$r(t) = \frac{r_0}{1 + \exp[-\gamma(x(t) + \tilde{\theta})]}, \quad (683)$$

where $\gamma \rightarrow \infty$ while $\langle x^2 \rangle$ stays constant. Is there a setting of the threshold θ which maximizes I_e ? Is there a cost to achieving this optimum?

One might worry that the high efficiency of coding seen in the fly’s H1 neuron arises because the fly has relatively few neurons, and thus is under greater pressure to be efficient. While this may be true, it seems that high coding efficiencies are there to be found even in animals like us and our primate cousins who have very large numbers of neurons. In humans it is possible to record from individual receptor cells in our hands and fingertips, contacting the axons of these cells as they course along the arm to the spinal cord. Data are more limited than in the fly, so one has to be more careful to avoid systematic errors, but the lower bound on the efficiency of coding complex, dynamic variations in the indentation of the skin is above 50%. In the visual cortex of non-human primates, there is a classic series of experiments correlating the perception of motion with the activity of single neurons in area MT. [probably this needs more explanation!] The standard stimuli for these experiments are random dot patterns in which a fraction of the dots move coherently while another fraction are randomly deleted and replaced at new locations; the perception of motion direction becomes less reliable as the degree of coherence decreases. The evidence that single neurons are making a measurable contribution to the perceptual decision is strong, since one can correlate the number of spikes generated by a neuron with the animal’s decision about leftward vs. rightward motion even when the coherence is zero, and the animal is just guessing.

The experiments in MT focused on asking the animal to report a decision about motion direction across a two second window of stimulation. When we look at these random patterns, however, we see a certain amount of “jiggling,” especially at low coherence. If we present exactly the same pattern of random dots vs. time, we find that the neurons respond with a fair degree of reliability to the temporal details of the movie, certainly down to time scales below 10 ms. In Fig 135 we see what this




FIG. 135 Entropy and information in spike trains from a motion sensitive neurons in the primate visual cortex (area MT); experiments by Britten et al (1993) and analysis by Strong et al (1998b). [fill in the caption]

means in terms of the information carried by the spike trains about the time-varying details of the visual stimulus, rather than just the overall direction of motion. Here the information, at $\Delta\tau = 6$ ms time resolution, is 25 – 30%. Experiments on the same neurons using stimuli that alternated between moving left and right⁸¹ on the 30 – 100 ms time scale found information rates of 1 – 2.5 bits/spike, quite comparable to the results with H1. In summary, although there are differences in the details of the spike trains from motion sensitive neurons in flies and monkeys, there is little different in the amount of information they carry, or the efficiency with this information is encoded, if we asking about the kinds of complex, dynamic stimuli that are relevant to the real world.

We now want to look at information transmission in the presence of noise, connecting back a bit to what we discussed in Chapters 1 and 2. Imagine that we are interested in some signal x , and we have a detector that generates data y which is linearly related to the signal

but corrupted by added noise:

$$y = gx + \xi. \quad (684)$$

It seems reasonable in many systems to assume that the noise is Gaussian, either for fundamental physical reasons (as with thermal noise), or because it arises from a superposition of many independent sources, in which case the central limit theorem takes over. We will also start with the assumption that x is drawn from a Gaussian distribution just because this is a simple place to start; we will see that we can use the maximum entropy property of Gaussians to make some more general statements based on this simple example. The question, then, is how much information observations on y provide about the signal x .

The problem of information transmission with Gaussian signals and noise is sufficiently important that it is worth going through all the algebra quite explicitly; this is also one of those pleasing problems where, as we calculate, terms proliferate and then collapse into a much simpler result. So, onward. The statement that ξ is Gaussian noise means that

$$P(y|x) = \frac{1}{\sqrt{2\pi\langle\xi^2\rangle}} \exp\left[-\frac{1}{2\langle\xi^2\rangle}(y-gx)^2\right]. \quad (685)$$

Our simplification is that the signal x also is drawn from a Gaussian distribution,

$$P(x) = \frac{1}{\sqrt{2\pi\langle x^2\rangle}} \exp\left[-\frac{1}{2\langle x^2\rangle}x^2\right], \quad (686)$$

and hence y itself is Gaussian,

$$P(y) = \frac{1}{\sqrt{2\pi\langle y^2\rangle}} \exp\left[-\frac{1}{2\langle y^2\rangle}y^2\right] \quad (687)$$

$$\langle y^2\rangle = g^2\langle x^2\rangle + \langle\xi^2\rangle. \quad (688)$$

To compute the information that y provides about x we use Eq. (672):

$$I(y \rightarrow x) = \int dy \int dx P(x, y) \log_2 \left[\frac{P(x, y)}{P(x)P(y)} \right] \quad \text{bits} \quad (689)$$

$$= \frac{1}{\ln 2} \int dy \int dx P(x, y) \ln \left[\frac{P(y|x)}{P(y)} \right] \quad (690)$$

$$= \frac{1}{\ln 2} \left\langle \ln \left[\frac{\sqrt{2\pi\langle y^2\rangle}}{\sqrt{2\pi\langle\xi^2\rangle}} - \frac{1}{2\langle\xi^2\rangle}(y-gx)^2 + \frac{1}{2\langle y^2\rangle}y^2 \right] \right\rangle, \quad (691)$$

where by $\langle \dots \rangle$ we understand an expectation value over

the joint distribution $P(x, y)$. Now in Eq. (691) we can

see that the first term is the expectation value of a constant. The third term involves the expectation value of y^2 divided by $\langle y^2 \rangle$, so we can cancel numerator and denominator. In the second term, we can take the expectation value first of y with x fixed, and then average over x , but since $y = gx + \xi$ the numerator is just the mean square fluctuation of y around its mean value, which again cancels with the $\langle \xi^2 \rangle$ in the denominator. So we have, putting the three terms together,

$$I(y \rightarrow x) = \frac{1}{\ln 2} \left[\ln \sqrt{\frac{\langle y^2 \rangle}{\langle \xi^2 \rangle}} - \frac{1}{2} + \frac{1}{2} \right] \quad (692)$$

$$= \frac{1}{2} \log_2 \left(\frac{\langle y^2 \rangle}{\langle \xi^2 \rangle} \right) \quad (693)$$

$$= \frac{1}{2} \log_2 \left(1 + \frac{g^2 \langle x^2 \rangle}{\langle \xi^2 \rangle} \right) \text{ bits.} \quad (694)$$

Another way of arriving at these results is to remember that the information is a difference of entropies [Eq (674)], but in this case the underlying distributions are all Gaussian. Thus it's useful to know, in general, the entropy of a Gaussian distribution. Suppose that

$$P(z) = \frac{1}{\sqrt{2\pi\langle(\delta z)^2\rangle}} \exp \left[-\frac{(z - \langle z \rangle)^2}{2\langle(\delta z)^2\rangle} \right]. \quad (695)$$

Now our task is to compute

$$S = - \int dz P(z) \log_2 P(z) = - \left\langle \log_2 P(z) \right\rangle. \quad (696)$$

But

$$\log_2 P(z) = \frac{1}{\ln 2} \left[\ln \left(\frac{1}{\sqrt{2\pi\langle(\delta z)^2\rangle}} \right) - \frac{(z - \langle z \rangle)^2}{2\langle(\delta z)^2\rangle} \right], \quad (697)$$

and hence

$$S = - \left\langle \log_2 P(z) \right\rangle \quad (698)$$

$$= \frac{1}{\ln 2} \left[\ln \left(\sqrt{2\pi\langle(\delta z)^2\rangle} \right) + \left\langle \frac{(z - \langle z \rangle)^2}{2\langle(\delta z)^2\rangle} \right\rangle \right] \quad (699)$$

$$= \frac{1}{\ln 2} \left[\frac{1}{2} \ln (2\pi\langle(\delta z)^2\rangle) + \frac{1}{2} \right] \quad (700)$$

$$= \frac{1}{2} \log_2 [2\pi e \langle(\delta z)^2\rangle]. \quad (701)$$

Notice that the entropy is independent of the mean, as we expect, since entropy measures variability or uncertainty.

Problem 138: Using the entropy of Gaussians. Use the general result on the entropy of Gaussian distributions, Eq (701),

to rederive Eq (694) for the information transmission through the “Gaussian channel.”

We can gain some intuition by rewriting Eq (694). Rather than thinking of our detector as adding noise after generating the signal gx , we can think of it as adding noise directly to the input, and then transducing this corrupted input:

$$y = g(x + \eta_{\text{eff}}), \quad (702)$$

where $\eta_{\text{eff}} = \xi/g$. Note that the “effective noise” η_{eff} is in the same units as the input x ; this is called “referring the noise to the input” and is a standard way of characterizing detectors, amplifiers and other devices, as discussed above.⁸² Written in terms of the effective noise level, the information transmission takes a simple form,

$$I(y \rightarrow x) = \frac{1}{2} \log_2 \left(1 + \frac{\langle x^2 \rangle}{\langle \eta_{\text{eff}}^2 \rangle} \right) \text{ bits,} \quad (703)$$

or

$$I(y \rightarrow x) = \frac{1}{2} \log_2 (1 + SNR), \quad (704)$$

where the signal to noise ratio is the ratio of the variance in the signal to the variance of the effective noise, $SNR = \langle x^2 \rangle / \langle \eta_{\text{eff}}^2 \rangle$.

The result in Eq. (704) is easy to picture: When we start, the signal is spread over a range $\delta x_0 \sim \langle x^2 \rangle^{1/2}$, but by observing the output of our detector we can localize the signal to a small range $\delta x_1 \sim \langle \eta_{\text{eff}}^2 \rangle^{1/2}$, and the reduction in entropy is $\sim \log_2(\delta x_0 / \delta x_1) \sim (1/2) \cdot \log_2(SNR)$, which is approximately the information gain.

Problem 139: A small point. Try to understand why the simple argument in the preceding paragraph, which seems sensible, doesn't give the exact answer for the information gain at small SNR .

⁸² As a reminder, if we build a photodetector it is not so useful to quote the noise level in Volts at the output—we want to know how this noise limits our ability to detect dim lights. Similarly, when we characterize a neuron that uses a stream of pulses to encode a continuous signal, we don't really want to know the variance in the pulse rate (although this is widely discussed); we want to know how noise in the neural response limits precision in estimating the real signal, and this amounts to defining an effective noise level in the units of the signal itself. In the present case this is just a matter of dividing, but generally it is a more complex task.




FIG. 136 Spatial profiles of Hunchback expression in the early *Drosophila* embryo. Small dots show experiments from individual embryos; circles with error bars are mean and standard deviation across 51 embryos. In the inset, image in red shows fluorescent antibody staining for Hb, and green shows the corresponding measurement for Krüppel. These images are taken by optical sectioning along the midline of the embryo, and the intensity is measured in a small area, roughly the size of a nucleus, the slides along the “rim” of the embryo where the nuclei are sitting. From Dubuis et al (2011).

To illustrate these ideas, consider the expression of the “gap genes” in the fly embryo, which we have seen in Sections [pointers to specific sections in previous chapters]. We recall that, in response to the primary, maternally supplied morphogens, these genes have varying levels of expression which provide a first step in building the blueprint for the fully developed organism. One of the basic ideas in developmental biology is that these expression levels carry “positional information,” i.e. that cells know where they are in the embryo, and hence their fate in the developed organism, as a result of knowing the concentrations of these molecules. It seems natural to ask if we can quantify this positional information, in

bits. To do this, as in Fig 136, we can look at many embryos and measure the concentration vs. position in each one. If there is a perfect functional relationship, with no noise, then the transmission of positional information is limited only by the number of samples that we take along the position axis, and hence the information in bits will just be the log of the number of cells. But there is noise, and this sets a limit to the positional information.

The position along the embryo can be measured by $0 \leq x \leq 1$. If we assume that the cells acquiring positional information are distributed uniformly (which is approximately true), then $P(x)$ is uniform, $P(x) = 1$. The expression level of the gene we are looking at will be called g . What we need to know is the distribution of expression levels at one position, $P(g|x)$. Experiments give us samples out of this distribution, but we may or may not have enough samples to characterize the whole distribution. What we can do more easily is to measure the mean $\bar{g}(x)$ and the variance $\sigma_g^2(x)$, and then approximate $P(g|x)$ as being Gaussian. One might worry that this approximation is uncontrolled, but in fact we can say more.

Suppose that all we know is the mean and variance of the distribution $P(g|x)$. The mutual information $I(g;x)$ is the difference between the entropy of the distribution $P(g)$ and the average entropy of the distribution $P(g|x)$,

$$I(g;x) = S[P(g)] - \langle S[P(g|x)] \rangle_x. \quad (705)$$

Thus if we can put an upper bound on the entropy $S[P(g|x)]$, we can put a lower bound on the information. Suppose we search for a distribution $P(g|x)$ that maximizes the entropy, while reproducing the measured mean and variance. As explained in more detail in Appendix [**], we can do this constrained optimization using the standard method of Lagrange multipliers. To maximize $S[P(g|x)]$ we introduce a functional

$$\tilde{S}[P(g|x)] = S[P(g|x)] - \lambda_1 \left[\int dg P(g|x)g - \bar{g}(x) \right] - \lambda_2 \left[\int dg P(g|x) (g - \bar{g}(x))^2 - \sigma_g^2(x) \right]. \quad (706)$$

Now if we maximize $\tilde{S}[P(g|x)]$ with respect to $P(g|x)$, and then extremize with respect to the Lagrange multipliers λ_1 and λ_2 , we will find a distribution that maximizes the entropy and reproduces the observed mean and variance. The solution to this problem, as shown in Appendix [**], is the Gaussian distribution. Thus, when we approximate $P(g|x)$ as being Gaussian, we generate a lower bound on the information $I(g;x)$.

In the example of Fig 136, this variance at each position is relatively small, with $\sigma_g(x) \sim 0.1$ in units where

the maximum mean expression level is one. Following through the computation of entropies as outlined above, one finds from these data that the expression level of Hunchback protein provides nearly two bits (give exact answer, with error bars) of information about position in the embryo. In the Gaussian approximation this is a lower bound on the information, but in fact the data sets are just large enough to make more direct estimates, and to show that this bound is tight [add a figure to illustrate this]. Classically, the gap genes have been described as

specifying boundaries, dividing the embryo into patches of high (on) and low (off) expression. Evidently a simple on/off picture corresponds at most to one bit of positional information, and so a quantitative analysis teaches us that the focus on “expression boundaries” literally misses half of the story.

Problem 140: Details of positional information. [Develop a problem that asks the students to use some of the real data on the gap genes ...]

As a next step consider the case where we observe several variables y_1, y_2, \dots, y_K in the hopes of

$$P(\{x_i\}) = \frac{1}{\sqrt{(2\pi)^K \det S}} \exp \left[-\frac{1}{2} x_i \cdot (S^{-1})_{ij} \cdot x_j \right] \quad (710)$$

$$P(\{y_i\}|\{x_i\}) = \frac{1}{\sqrt{(2\pi)^K \det N}} \exp \left[-\frac{1}{2} (y_j - g_{ik}x_k) \cdot (N^{-1})_{ij} \cdot (y_j - g_{jm}x_m) \right], \quad (711)$$

where again the summation convention is used; $\det S$ denotes the determinant of the matrix S , $(S^{-1})_{ij}$ is the ij element in the inverse of the matrix S , and similarly for the matrix N .

To compute the mutual information we proceed as before. First we find $P(\{y_i\})$ by doing the integrals over the x_i ,

$$P(\{y_i\}) = \int d^K x P(\{y_i\}|\{x_i\})P(\{x_i\}), \quad (712)$$

and then we write the information as an expectation value,

$$I(\{y_i\} \rightarrow \{x_i\}) = \left\langle \log_2 \left[\frac{P(\{y_i\}|\{x_i\})}{P(\{y_i\})} \right] \right\rangle, \quad (713)$$

where $\langle \dots \rangle$ denotes an average over the joint distribution $P(\{y_i\}, \{x_i\})$. As in Eq. (691), the logarithm can be broken into several terms such that the expectation value of each one is relatively easy to calculate. Two of three terms cancel, and the one which survives is related to the normalization factors that come in front of the exponentials. After the dust settles we find

$$I(\{y_i\} \rightarrow \{x_i\}) = \frac{1}{2} \text{Tr} \log_2 [\mathbf{1} + N^{-1} \cdot (g \cdot S \cdot g^T)], \quad (714)$$

where Tr denotes the trace of a matrix, $\mathbf{1}$ is the unit matrix, and g^T is the transpose of the matrix g .

learning about the same number of underlying signals x_1, x_2, \dots, x_K . The equations analogous to Eq. (684) are then

$$y_i = g_{ij}x_j + \xi_i, \quad (707)$$

with the usual convention that we sum over repeated indices. The Gaussian assumptions are that each x_i and ξ_i has zero mean, but in general we have to think about arbitrary covariance matrices,

$$S_{ij} = \langle x_i x_j \rangle \quad (708)$$

$$N_{ij} = \langle \xi_i \xi_j \rangle. \quad (709)$$

The relevant probability distributions are

Problem 141: The multi-dimensional Gaussian. Fill in the details leading to Eq (714). [where do I give the problem $\text{Tr} \ln = \ln \det$? [connect here](#)]

The matrix $g \cdot S \cdot g^T$ describes the covariance of those components of y that are contributed by the signal x . We can always rotate our coordinate system on the space of y s to make this matrix diagonal, which corresponds to finding the eigenvectors and eigenvalues of the covariance matrix; these eigenvectors are also called “principal components.” For a Gaussian distribution, the eigenvectors describe directions in the space of y which are fluctuating independently, and the eigenvalues are the variances along each of these directions. If the covariance of the noise is diagonal in the same coordinate system, then the matrix $N^{-1} \cdot (g \cdot S \cdot g^T)$ is diagonal and the elements along the diagonal are the signal to noise ratios along each independent direction. Taking the $\text{Tr} \log$ is equivalent to computing the information transmission along each direction using Eq. (704), and then summing the results.

An important case is when the different variables x_i represent a signal sampled at several different points in

time. Then there is some underlying continuous function $x(t)$, and in place of the discrete Eq. (707) we have the continuous linear response of the detector to input signals,

$$y(t) = \int dt' M(t-t')x(t') + \xi(t). \quad (715)$$

In this continuous case the analog of the covariance matrix $\langle x_i x_j \rangle$ is the correlation function $\langle x(t)x(t') \rangle$. We are usually interested in signals (and noise) that are stationary. This means, as discussed in Appendix A.2, that all statistical properties of the signal are invariant to translations in time: a particular pattern of wiggles in the function $x(t)$ is equally likely to occur at any time. Thus, the correlation function which could in principle depend on two times t and t' depends only on the time difference,

$$\langle x(t)x(t') \rangle = C_x(t-t'). \quad (716)$$

The correlation function generalizes the covariance matrix to continuous time, but we have seen that it can be useful to diagonalize the covariance matrix, thus finding a coordinate system in which fluctuations in the different directions are independent. From [pointer] we know that the answer is to go into a Fourier representation, where (in the Gaussian case) different Fourier components are independent and their variances are (up to normalization) the power spectra.

To complete the analysis of the continuous time Gaussian channel described by Eq. (715), we again refer noise

to the input by writing

$$y(t) = \int dt' M(t-t')[x(t') + \eta_{\text{eff}}(t')]. \quad (717)$$

If both signal and effective noise are stationary, then each has a power spectrum; let us denote the power spectrum of the effective noise η_{eff} by $N_{\text{eff}}(\omega)$ and the power spectrum of x by $S_x(\omega)$ as usual. There is a signal to noise ratio at each frequency,

$$SNR(\omega) = \frac{S_x(\omega)}{N_{\text{eff}}(\omega)}, \quad (718)$$

and since we have diagonalized the problem by Fourier transforming, we can compute the information just by adding the contributions from each frequency component, so that

$$I[y(t) \rightarrow x(t)] = \frac{1}{2} \sum_{\omega} \log_2[1 + SNR(\omega)]. \quad (719)$$

Finally, to compute the frequency sum, we recall that [I think this is found also in an Appendix; check!]

$$\sum_n f(\omega_n) \rightarrow T \int \frac{d\omega}{2\pi} f(\omega). \quad (720)$$

Thus, the information conveyed by observations on a (large) window of time becomes

$$I[y(0 < t < T) \rightarrow x(0 < t < T)] \rightarrow \frac{T}{2} \int \frac{d\omega}{2\pi} \log_2[1 + SNR(\omega)] \text{ bits}. \quad (721)$$

We see that the information gained is proportional to the time of our observations, so it makes sense to define an information rate:

$$\begin{aligned} R_{\text{info}} &\equiv \lim_{T \rightarrow \infty} \frac{1}{T} \cdot I[y(0 < t < T) \rightarrow x(0 < t < T)] \\ &= \frac{1}{2} \int \frac{d\omega}{2\pi} \log_2[1 + SNR(\omega)] \text{ bits/sec.} \end{aligned} \quad (722)$$

Note that in all these equations, integrals over frequency run over both positive and negative frequencies; if the signals are sampled at points in time spaced by τ_0 then the maximum (Nyquist) frequency is $|\omega|_{\text{max}} = \pi/\tau_0$.

Problem 142: How long to look? We know that when we integrate for longer times we can suppress the effects of noise and

hence presumably gain more information. Usually we would say that the benefits of integration are cut off by the fact that the signals we are looking at will change. But once we think about information transmission there is another possibility—perhaps we would learn more by using the same time to look at something new, rather than getting a more accurate view of something we have already seen. To address this possibility, let's consider the following simple model. We look at one thing for a time τ , and then jump to something completely new. Given that we integrate for τ , we achieve some signal-to-noise ratio which we'll call $S(\tau)$.

(a.) Explain why, in this simple model, if the noise is Gaussian then the rate at which we gain information is at most

$$R_{\text{info}}(\tau) = \frac{1}{\tau} \log_2[1 + S(\tau)]. \quad (724)$$

How does the assumption that we 'jump to something completely new' enter into the justification of this formula?

(b.) To make progress we need a model for $S(\tau)$. Since this is the signal-to-noise ratio let's start with the signal. Suppose that inputs are given by x , and the output is y . At $t = 0$, the value of y is set to zero, and after that our sensory receptor responds to its inputs according to a simple differential equation

$$\tau_0 \frac{dy}{dt} = -y + x. \quad (725)$$

Show that $y(\tau) = x[1 - \exp(-\tau/\tau_0)]$. Now for the noise, suppose that $\eta_{\text{eff}}(t)$ has a correlation function

$$\langle \eta_{\text{eff}}(t)\eta_{\text{eff}}(t') \rangle = \sigma_0^2 e^{-|t-t'|/\tau_c}. \quad (726)$$

Show that if we average the noise over a window of duration τ , then the variance

$$\sigma^2(\tau) \equiv \left\langle \left[\frac{1}{\tau} \int_0^\tau dt \eta_{\text{eff}}(t) \right]^2 \right\rangle \approx \sigma_0^2 \quad (\tau \ll \tau_0) \quad (727)$$

$$\approx \frac{2\sigma_0^2\tau_c}{\tau} \quad (\tau \gg \tau_0). \quad (728)$$

Give a more general analytic expression for $\sigma^2(\tau)$. Put these factors together to get an expression for $S(\tau) = y^2(\tau)/\sigma^2(\tau)$. To keep things simple, you can assume that the time scale which determines the response to inputs is the same as that which determines the correlations in the noise, so that $\tau_c = \tau_0$.

(c.) Hopefully you can show from your results in [b] that $S(\tau \gg \tau_0) \propto \tau$. This corresponds to our intuition that signal-to-noise ratios grow with averaging time because we beat down the noise, not worrying about the possibility that the signal itself will change. What happens for $\tau \ll \tau_0$?

(d.) Suppose that τ_0 is very small, so that all “reasonable” values of $\tau \gg \tau_0$. Then, from [c], $S(\tau) = A\tau$, with A a constant. With this assumption, plot $R_{\text{info}}(\tau)$; show that with proper choice of units, you don’t need to know the value of A . What value of τ maximizes the information rate? Is this consistent with the assumption that $\tau \gg \tau_0$?

(e.) In general, the maximum information is found at the point where $dR_{\text{info}}/d\tau = 0$. Show that this condition can be rewritten as a relationship between the signal-to-noise ratio and its logarithmic derivative, $z = d \ln S(\tau)/d \ln \tau$. From your previous results, what can you say about the possible values of z as τ is varied? Use this to bound $S(\tau)$ at the point of maximum R_{info} . What does this say about the compromise between looking carefully at one thing and jumping to something new?

(f.) How general can you make the conclusions that you draw in [e]?

In the same way that we used the Gaussian approximation to put bounds on the positional information carried by the gap genes, we can put bounds on the information carried by sensory neurons. As discussed in Section [**], we can reconstruct continuous sensory input signals from the discrete sequences of action potentials, sometimes quite accurately. Concretely, the sensory stimulus $s(t)$ could be light intensity as a function of time in a small region of the visual field, sound pressure as a function of time at the ear canal, the amplitude of mechanical vibrations in sensors such as the cricket cercus and frog sacculus, We can estimate the signal from the spike times $\{t_i\}$ in a single neuron as

$$s_{\text{est}}(t) = \sum_i f(t - t_i), \quad (729)$$

where the filter $f(\tau)$ is chosen to minimize $\chi^2 = \langle |s_{\text{est}}(t) - s(t)|^2 \rangle$. Then the quality of the reconstructions can be evaluated by measuring the power spectrum of errors in the reconstruction, and referring these errors to the input, frequency component by frequency component,

$$\tilde{s}_{\text{est}}(\omega) = g(\omega) [\tilde{s}(\omega) + \tilde{\eta}_{\text{eff}}(\omega)]. \quad (730)$$




FIG. 137 Coding efficiency in cricket and frog vibration sensors. (A) A schematic of experiments on the cricket cercal sensors, with direct stimulation of the sensory hairs and recording from the primary sensory neurons. (B) Stimulus (dashed) and reconstruction (Solid line) in experiments on the cercal neurons. (C) Power spectral density of the signal, and the noise η_{eff} in the reconstructions, from Eq(730). (D) Coding efficiency for example neurons in the cricker cercus and the frog sacculus, using successively higher order approximations to the spike train entropy. Variable timing precision is implemented by providing the reconstruction algorithm in Eq (729) with spike times t_i at limited resolution. From Rieke et al (1993).

Although the errors in the reconstruction might not be exactly Gaussian, the maximum entropy argument above tells us that we can put a lower bound on the information which the spike train provides about the stimulus $s(t)$ by measuring the power spectrum of the effective noise η_{eff} . An example is shown in Fig 137, from experiments on the mechanical sensors in the cricket and frog. Importantly, we can also put upper bounds on the entropy of the spike train, first by assuming that spikes occur independently, then by assuming that the intervals between spikes are independent, then allowing for correlations between successive intervals. With a lower bound on the information and an upper bound on the entropy, we have a lower bound on their ratio, the coding efficiency. In these systems, as with the case of H1 in Fig 134, we see that efficiencies reach $\sim 50\%$ with timing precision in the millisecond range.

By now both the “direct” and the “reconstruction” methods have been used to measure information rates and coding efficiencies in a wide range of neurons responding to sensory stimuli, from the first steps of sensory coding in invertebrates, such as the cricket cercal system in Fig 137, to cells deep in primate visual cortex. The result that single neurons use 30–50% of their spike train entropy to encode sensory information, even down to millisecond resolution, has been confirmed in many systems [maybe reminder that references are at the end of the section?]. An important thread running through

this work is that information rates and coding efficiencies are higher, and the high coding efficiency extends to higher time resolution, when sensory inputs are more like those which occur in nature—complex, dynamic, and with enormous dynamic range; an example from the frog auditory system is shown in Fig 138 [do we need more examples here?]. These results suggest not only that the brain is capable of efficient coding, but also that this efficiency is achieved by matching neural coding strategies to the structure of natural sensory inputs. We will return to this idea in Section IV.C.

The Gaussian channel gives us the opportunity to explore the way in which noise limits information transmission. Imagine that we have measured the spectrum of the effective noise, $N_{\text{eff}}(\omega)$. By changing the spectrum of input signals, $S(\omega)$, we can change the rate of information transmission. Can we maximize this information rate? Clearly this problem is not well posed without some constraints: if we are allowed just to increase the amplitude of the signal—multiply the spectrum by a large constant—then we can always increase information transmission. We need to study the optimization of information rate with some fixed ‘dynamic range’ for the signals. A simple example, considered by Shannon at the outset, is to fix the total variance of the signal, which is the same as fixing the integral of the spectrum. We can motivate this constraint by noting that if the signal is a voltage and we have to drive this signal through a re-




FIG. 138 Coding efficiency in frog auditory neurons. At left, the power spectrum of a broadband, artificial stimulus (top) and a stimulus shaped to have the same spectrum as bullfrog calls (bottom). These stimuli were played to the bullfrog while recording from individual auditory neurons emerging from the amphibian papilla. Reconstructing the sound pressure as a function of time allows us to bound the information transmission rate, as explained in the text, and from this we estimate the coding efficiency—the ratio of the information rate to the entropy rate. In this example, at right, we see clearly that the coding efficiency is substantially higher for the more naturalistic stimuli, approaching 90%. From Rieke et al (1995).

sistive element, then the variance is proportional to the mean power dissipation. Alternatively, it might be easy to measure the variance of the signals that we are interested in (as for the visual signals in the example below), and then the constraint is empirical.

So the problem we want to solve is maximizing R_{info} while holding $\langle x^2 \rangle$ fixed. As before, we introduce a Lagrange multiplier and maximize a new function

$$\tilde{R} = R_{\text{info}} - \lambda \langle x^2 \rangle \quad (731)$$

$$= \frac{1}{2} \int \frac{d\omega}{2\pi} \log_2 \left[1 + \frac{S_x(\omega)}{N_{\text{eff}}(\omega)} \right] - \lambda \int \frac{d\omega}{2\pi} S_x(\omega). \quad (732)$$

The value of the function $S_x(\omega)$ at each frequency contributes independently, so it is easy to compute the functional derivatives,

$$\frac{\delta \tilde{R}}{\delta S_x(\omega)} = \frac{1}{2 \ln 2} \cdot \frac{1}{1 + S_x(\omega)/N_{\text{eff}}(\omega)} \cdot \frac{1}{N_{\text{eff}}(\omega)} - \lambda, \quad (733)$$

and the optimization condition is $\delta \tilde{R}/\delta S_x(\omega) = 0$. The result is that

$$S_x(\omega) + N_{\text{eff}}(\omega) = \frac{1}{2\lambda \ln 2}. \quad (734)$$

Thus the optimal choice of the signal spectrum is one which makes the sum of signal and (effective) noise equal to white noise! This, like the fact that information is maximized by a Gaussian signal, is telling us that efficient information transmission occurs when the received signals are as random as possible given the constraints. Thus an attempt to look for structure in an optimally encoded signal (say, deep in the brain) will be frustrating.

In general, complete whitening as suggested by Eq. (734) can't be achieved at all frequencies, since if the system has finite time resolution (for example) the effective noise grows without bound at high frequencies. Thus the full solution is to have the spectrum determined by Eq. (734) everywhere that the spectrum comes out to a positive number, and then to set the spectrum equal to zero outside this range. If we think of the effective noise spectrum as a landscape with valleys, the condition for optimizing information transmission corresponds to filling the valleys with water; the total volume of water is the variance of the signal.

Problem 143: Whitening. Consider a system that responds linearly to a signal $s(t)$, with added noise $\eta(t)$:

$$x(t) = \int d\tau F(\tau) s(t - \tau) + \eta(t). \quad (735)$$

Assume that the noise is Gaussian and white, with power spectrum \mathcal{N}_0 , so that

$$\langle \eta(t)\eta(t') \rangle = \mathcal{N}_0 \delta(t - t'). \quad (736)$$

For simplicity, assume that the signal $s(t)$ is Gaussian, with a power spectrum $S(\omega)$,

$$\langle s(t)s(t') \rangle = \int \frac{d\omega}{2\pi} S(\omega) \exp[-i\omega(t-t')]. \quad (737)$$

(a.) Write an expression for the rate R_{info} at which the observable $x(t)$ provides information about the signal $s(t)$.

(b.) The variance of the variable $x(t)$ is not well defined. Why? Consider just the component of $x(t)$ that comes from the signal $s(t)$, that is Eq (735) but with $\eta = 0$. Find an expression for the variance of this “output signal.”

(c.) Consider the problem of maximizing R_{info} by adjusting the filter $F(\tau)$. Obviously the information transmission is larger if F is larger, so to make the problem well posed assume that the variance of the output signal (from [b]) is fixed. Show that this variational problem can be solved explicitly for $|\hat{F}(\omega)|^2$, where $\hat{F}(\omega)$ is the Fourier transform of the filter $F(\tau)$. Can you explain intuitively why only the modulus, and not the phase, of $\hat{F}(\omega)$ is relevant here?

(d.) Find the limiting form of the optimal filter as the noise becomes small. What does this filter do to the input signal? Explain why this makes sense. Saying that “noise is small” is slightly strange, since \mathcal{N}_0 has units. Give a more precise criterion for your small noise limit to be valid.

(e.) Consider the case of an input with exponentially decaying correlations, so that

$$S(\omega) = \frac{2\langle s^2 \rangle \tau_c}{1 + (\omega\tau_c)^2}, \quad (738)$$

where τ_c is the correlation time. Find the optimal filter in this case, and use this to evaluate the maximum value of R_{info} as a function of the output signal variance. You should check that your results for R_{info} , which should be in bits/s, are independent of the units used for the output variance and the noise power spectrum. Contrast your result with what would happen if $|\hat{F}(\omega)|$ were flat as a function of frequency, so that there was no real filtering (just a multiplication so that the output signal variance comes out right). How much can one gain by building the right filter?

These ideas have been used to characterize information transmission across the first synapse in the fly’s visual system. We have seen these data before, in thinking about how the precision of photon counting changes as the background light intensity increases. Recall from Section I.A that, over a reasonable dynamic range of intensity variations, the average voltage response of the photoreceptor cell is related linearly to the intensity or contrast in the movie, and the noise or variability $\delta V(t)$ is governed by a Gaussian distribution of voltage fluctuations around the average:

$$V(t) = V_{\text{DC}} + \int dt' T(t-t') C(t') + \delta V(t). \quad (739)$$

This (happily) is the problem we have just analyzed.

As before, we think of the noise in the response as being equivalent to noise $\delta C_{\text{eff}}(t)$ that is added to the movie itself,

$$V(t) = V_{\text{DC}} + \int dt' T(t-t') [C(t') + \delta C_{\text{eff}}(t')]. \quad (740)$$

Since the fluctuations have a Gaussian distribution, they can be characterized completely by their power spectrum

$N_C^{\text{eff}}(\omega)$, which measures the variance of the fluctuations that occur at different frequencies,

$$\langle \delta C_{\text{eff}}(t) \delta C_{\text{eff}}(t') \rangle = \int \frac{d\omega}{2\pi} N_C^{\text{eff}}(\omega) \exp[-i\omega(t-t')]. \quad (741)$$

There is a minimum level of this effective noise set by the random arrival of photons (shot noise). The photon noise is white if expressed as $N_C^{\text{eff}}(\omega)$, although it makes a nonwhite contribution to the voltage noise. As we have discussed, over a wide range of background light intensities and frequencies, the fly photoreceptors have effective noise levels that reach the limit set by photon statistics. At high frequencies there is excess noise beyond the physical limit, and this excess noise sets the time resolution of the system.

The power spectrum of the effective noise tells us, ultimately, what signals the photoreceptor can and cannot transmit. How do we turn these measurements into bits? One approach is to assume that the fly lives in some particular environment, and then calculate how much information the receptor cell can provide about this particular environment. But to characterize the cell itself, we might ask a different question: in principle how much information can the cell transmit? To answer this question we are allowed to shape the statistical structure of the environment so as to make the best use of the receptor (the opposite, presumably, of what happens in evolution!). This is just the optimization discussed above, so it is possible to turn the measurements on signals and noise into estimates of the information capacity of these cells. This was done both for the photoreceptor cells and for the large monopolar cells (LMCs) that receive direct synaptic input from a group of six receptors. From measurements on natural scenes the mean square contrast signal was fixed at $\langle C^2 \rangle = 0.1$. Results are shown in Fig 139.

The first interesting feature of the results is the scale: individual neurons are capable of transmitting well above 1000 bits per second. This does not mean that this capacity is used under natural conditions, but rather speaks to the precision of the mechanisms underlying the detection and transmission of signals in this system. Second, information capacity continues to increase as the level of background light increases: noise due to photon statistics is less important in brighter lights, and this reduction of the physical limit actually improves the performance of the system even up to very high photon counting rates, indicating once more that the physical limit is relevant to the real performance. Third, we see that the information capacity as a function of photon counting rate is shifted along the counting rate axis as we go from photoreceptors to the LMCs, and this corresponds (quite accurately!) to the fact that LMCs integrate signals from six photoreceptors and thus act as if they captured photons at a six times higher rate. Finally, in the large monopolar cells in-




FIG. 139 At left, the effective contrast noise levels in a single photoreceptor cell, a single LMC (the second order cell) and the inferred noise level for a single active zone of the synapse from photoreceptor to LMC. The hatching shows the signal spectra required to whiten the total output over the largest possible range while maintaining the input contrast variance $\langle C^2 \rangle = 0.1$, as discussed in the text. At right, the resulting information capacities as a function of the photon counting rates in the photoreceptors. From de Ruyter van Steveninck & Laughlin (1996).

formation has been transmitted across a synapse, and in the process is converted from a continuous voltage signal into discrete events corresponding to the release of neurotransmitter vesicles at the synapse. As a result, there is a new limit to information transmission that comes from viewing the large monopolar cell as a “vesicle counter.”

[This discussion needs to be fleshed out. It's also the second independent use of max ent in this section, which makes me worry that leaving max ent to an Appendix may be a mistake, although it also comes up earlier .. this is a pretty big organizational issue. Also was thinking of being explicit about max ent for counting, above, which would make things easier here! If every vesicle makes a measurable, deterministic contribution to the cell's response (a generous assumption), then the large monopolar cell's response is equivalent to reporting how many vesicles are counted in a small window of time corresponding to the photoreceptor time resolution. We don't know the distribution of these counts, but we can estimate (from other experiments, with uncertainty) the mean count, and we know that there is a maximum entropy for any count distribution once we fix the mean (see, for example, Appendix A.8). No mechanism at the synapse can transmit more information than this limit. Remarkably, the fly operates within a factor of two of this limit, and the agreement might be even better but for uncertainties in the vesicle counting rate.

[This section needs a summary and conclusion!]

To a remarkable extent, Shannon's original work provides a complete and accessible guide to the foundations of the subject (Shannon 1948). Seldom has something genuinely new emerged so fully in one (admittedly long, two part) paper. For a modern textbook account, the standard is set by Cover & Thomas (1991). An fascinating if idiosyncratic treatment of Shannon's ideas is given by Brillouin (1962). A recent textbook that emphasizes connections between information theory and statistical physics is Mézard & Montanari (2009). The brief discussion of four letter words is based on Stephens & Bialek (2010).

Brillouin 1962: *Science and Information Theory* L Brillouin (Academic, New York, 1962).

Cover & Thomas 1991: *Elements of Information Theory* TM Cover & JA Thomas (Wiley, New York, 1991); there is also a second edition (2006).

Mézard & Montanari 2009: *Information, Physics and Computation*. M Mézard & A Montanari (Oxford University Press, Oxford, 2009).

Shannon 1948: A mathematical theory of communication, CE Shannon, *Bell Sys. Tech. J.* **27**, 379–423 & 623–656 (1948). Reprinted in CE Shannon & W Weaver, *The Mathematical Theory of Communication* (University of Illinois Press, Urbana, 1949).

Stephens & Bialek 2010: Statistical mechanics of letters in words. GJ Stephens & W Bialek, *Phys Rev E* **81**, 066119 (2010); arXiv:0801.0253 [q-bio.NC] (2008).

The exploration of neural coding using ideas from information theory rests on a large literature, starting with Adrian's first experiments recording the spikes from individual sensory neurons in the 1920s. For a guide to the field up to the mid 1990s, see Rieke et al (1997). The idea of using ergodicity to make “direct” estimates of the entropy and information in spike trains as they encode dynamic signals is presented by Strong et al (1998a) and de Ruyter van Steveninck et al (1997). For the original discussion of the limits to information transmission by spike trains, see MacKay & McCulloch (1952). For the analysis of information carried by single events, see DeWeese & Meister (1999) and Brenner et al (2000).

Brenner et al 2000: Synergy in a neural code. N Brenner, SP Strong, R Koberle, W Bialek & RR de Ruyter van Steveninck, *Neural Comp* **12**, 1531–1552 (2000); arXiv:physics/9902067 (1999).

DeWeese & Meister 1999: How to measure the information gained from one symbol. MR DeWeese & M Meister, *Network* **10**, 325–340 (1999).

MacKay & McCulloch 1952: The limiting information capacity of a neuronal link. D MacKay & WS McCulloch, *Bull Math Biophys* **14**, 127–135 (1952).

de Ruyter van Steveninck et al 1997: Reproducibility and variability in neural spike trains. RR de Ruyter van Steveninck, GD Lewen, SP Strong, R Koberle & W Bialek, *Science* **275**, 1805–1808 (1997).

Strong et al 1998a: Entropy and information in neural spike trains. SP Strong, R Koberle, RR de Ruyter van Steveninck & W Bialek, *Phys Rev Lett* **80**, 197–200 (1998); arXiv:cond-mat/9603127 (1996).

The experiments connecting motion perception to the activity of individual neurons in visual cortex were reported in a series of beautiful papers by Newsome, Movshon, and their collaborators (Newsome et al 1989) [give a proper guide!]. Of particular relevance here is Britten et al (1993). The fact that these experiments generated demonstrated reproducible responses to stimuli that repeated their temporal details was emphasized by Bair & Koch (1996). The

analysis in Fig 135 is from Strong et al (1998b), unpacking a footnote in Strong et al (1998a) above. Experiments designed to look more specifically at these problems of information transmission for dynamic signals in MT were done by Buračas et al (1998).

Bair & Koch 1996: Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey, W Bair & C Koch, *Neural comp* **8**, 1185–1192 (1996).

Britten et al 1993: Response of neurons in macaque MT to stochastic motion signals. KH Britten, MN Shadlen, WT Newsome & JA Movshon, *Vis Neuroci* **10**, 1157–1169 (1993).

Buračas et al 1998: Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. GT Buračas, AM Zador, MR DeWeese & TD Albright, *Neuron* **20**, 959–969 (1998).

Newsome et al 1989: Neuronal correlates of a perceptual decision. WT Newsome, KH Britten & JA Movshon, *Nature* **341**, 52–54 (1989).

Strong et al 1998b: On the application of information theory to neural spike trains. SP Strong, RR de Ruyter van Steveninck, W Bialek & R Koberle, in *Pacific Symposium on Biocomputing '98*, RB Altman, AK Dunker, L Hunter & TE Klein, eds, pp 621–632 (World Scientific, Singapore, 1998).

The classic discussion of information transmission in the presence of Gaussian noise is again by Shannon (1949), and again the standard modern textbook account is in Cover and Thomas (1991), cited in Section IV.A. The discussion of positional information is based on Dubuis et al (2011). The idea of decoding the spike train to recover the underlying signal (“stimulus reconstruction”) was introduced in Section [pointer], and is reviewed by Rieke et al (1997), and the first use of this approach to compare information and entropy was by Rieke et al (1993). This led to experiments showing that coding efficiencies are higher for more naturalistic stimuli (Rieke et al 1995), and this problem was eventually revisited in the context of much more natural stimuli using the “direct” methods of information estimation coupled with more sophisticated strategies for dealing with finite sampling Nemenman et al (2008), as explained in Appendix A.9. The measurements on information capacity in the fly retina are by de Ruyter van Steveninck and Laughlin (1996). The comparison of this information rate with the limits set by counting vesicles is discussed by Rieke et al (1997), above [probably need to unpack the original refs about vesicle counting rates].

Dubuis et al 2011: Positional information, in bits. JO Dubuis, G Tkačik, W Bialek, EF Wieschaus & T Gregor, in preparation (2011).

Nemenman et al 2008: Neural coding of a natural stimulus ensemble: Information at sub-millisecond resolution. I Nemenman, GD Lewen, W Bialek & RR de Ruyter van Steveninck, *PLoS Comp Bio* **4**, e1000025 (2008); arXiv:q-bio.NC/0612050 (2006).

Rieke et al 1993: Coding efficiency and information rates in sensory neurons. F Rieke, D Warland & W Bialek, *Europhys. Lett* **22**, 151–156 (1993).

Rieke et al 1995: Naturalistic stimuli increase the rate and efficiency of information transmission by primary auditory neurons. F Rieke, DA Bodnar & W Bialek, *Proc R. Soc Lond. Ser. B* **262**, 259–265 (1995).

Rieke et al 1997: *Spikes: Exploring the Neural Code* F Rieke, D Warland, RR de Ruyter van Steveninck & W Bialek (MIT Press, Cambridge, 1997).

de Ruyter van Steveninck & Laughlin 1996: The rate of information transfer at graded-potential synapses. RR de Ruyter van Steveninck & SB Laughlin, *Nature* **379**, 642–645 (1996).

Shannon 1949: Communication in the presence of noise. CE Shannon, *Proc IRE* **37**, 10–21 (1949).

I need to add a guide to all the experiments on information rates and coding efficiencies, using both direct and reconstruction methods. Even this list is incomplete.

Attias & Schreiner 1998: Coding of naturalistic stimuli by auditory midbrain neurons. H Attias & CE Schreiner, in *Advances in Neural Information Processing Systems 10*, MI Jordan, MJ Kearns & SA Solla, pp 103–109 (MIT Press, Cambridge, 1998).

Bair & Koch 1996: Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. W Bair & C Koch, *Neural Comp* **8**, 44–66 (1996).

Buračas et al 1998: Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. GT Buračas, AM Zador, MR DeWeese & TD Albright, *Neuron* **20**, 959–969 (1998).

Escabi et al 2003: Naturalistic auditory contrast improves spectrotemporal coding in the cat inferior colliculus. MA Escabi, LM Miller, HL Read & CE Schreiner, *J Neurosci* **23**, 11489–11504 (2003).

Kara et al 2000: Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. P Kara, P Reinagel & RC Reid, *Neuron* **27**, 635–646 (2000).

Koch et al 2006: How much the eye tells the brain. K Koch, J McLean, R Segev, MA Freed, MJ Berry II, V Balasubramanian & P Sterling, *Curr Biol* **16**, 1428–1434 (2006).

Liu et al 2001: Variability and information in a neural code of the cat lateral geniculate nucleus. RC Liu, S Tzonev, S Rebrik & KD Miller, *J Neurophysiol* **86**, 2789–2806 (2001).

Reinagel & Reid 2000: Temporal coding of visual information in the thalamus. P Reinagel & RC Reid, *J Neurosci* **20**, 5392–5400 (2000).

Rokem et al 2006: Spike-timing precision underlies the coding efficiency of auditory receptor neurons. A Rokem, S Watzl, T Gollisch, M Stemmler, AVM Herz & I Samengo, *J Neurophysiol* **95**, 2541–2552 (2006).

Simmons & de Ruyter van Steveninck 2010: Sparse but specific temporal coding by spikes in an insect sensory-motor ocellar pathway. PJ Simmons & RR de Ruyter van Steveninck, *J Exp Biol* **213**, 2629–2639 (2010).

Yu et al 2005: Preference of sensory neural coding for $1/f$ signals. Y Yu, R Romero & TS Lee, *Phys Rev Lett* **94**, 108103 (2005).

I haven't said anything about error correcting codes. I don't see, in the short run, how to connect these elegant ideas to real biological phenomena. On the other hand, they are so interesting ... at the very least I will need to give references, and some commentary about why we should be trying to think about this.

B. Does biology care about bits?

The question for this section has been with us almost since Shannon's original work. On the one hand, the few examples we have seen in the last section certainly suggest that organisms are squeezing more bits out of

their hardware than we might naively have expected, perhaps even coming close to physical limits on information transmission. On the other hand, the usual view of information theory is as a theory for communication, with its most sophisticated developments in the context of error correcting codes, which seem of little relevance to the natural (as opposed to the engineered) world. Here we'll review old ideas about the connection of information to gambling, and see how closely related ideas have reappeared in thinking about the life strategies of bacterial populations. Then we'll step back and try to look more generally at the connections among information, biological function and evolutionary fitness, and argue that evolution really can select for biological mechanisms that are efficient in an information theoretic sense.

To start, let us consider a simple game; this may seem like a strange topic for a physics course, but please bear with me! I will flip a coin, and you bet on whether it will come up heads or tails. If you get it right, I double your money. If you're wrong, you lose what you bet. If this is a fair coin, so that heads and tails each come up half the time, there really isn't anything to analyze, what happens is "just chance." But if you know, for example, that this is a biased coin, and that the probability of heads really is 60%, you might be tempted to put all of your money on heads. On average, if you bet one dollar you will receive $2 \times (0.6) = 1.2$ dollars in return, which sounds good. Indeed, if we play only once then this is what you should do, since it will maximize your expected return.

But what happens if we are going to play repeatedly, which you might think is a better metaphor for life? Now if you put all your money on heads, there is a 40% chance that, in one flip, you'll lose it all. Suppose that instead you put a fraction f of your money on heads and a fraction $1-f$ on tails. If we introduce a binary variable $n = 1$ for heads and $n = 0$ for tails, then on the i^{th} flip your winnings will change by a factor

$$G_i = 2 \times [fn_i + (1-f)(1-n_i)], \quad (742)$$

where n_i marks what happens on the i^{th} flip. After N successive flips you will have a gain

$$G_{\text{total}}(N) = 2^N \prod_{i=1}^N [fn_i + (1-f)(1-n_i)], \quad (743)$$

where we are assuming that you consistently put a fraction f of your accumulated winnings down as a bet on

heads, and the remainder on tails.

To keep going, we want to write the product in Eq (743) as the exponential of a sum. It's useful to notice that, because n_i is either 0 or 1, we have

$$fn_i + (1-f)(1-n_i) = \exp[n_i \ln(f) + (1-n_i) \ln(1-f)]. \quad (744)$$

This means that we can write the total gain

$$\begin{aligned} G_{\text{total}}(N) &= 2^N \prod_{i=1}^N [fn_i + (1-f)(1-n_i)] \\ &= 2^N \prod_{i=1}^N \exp[n_i \ln(f) + (1-n_i) \ln(1-f)] \end{aligned} \quad (745)$$

$$= \exp[N\Lambda(f; \{n_i\})], \quad (746)$$

where

$$\Lambda(f; \{n_i\}) = \ln 2 + \frac{1}{N} \sum_{i=1}^N [n_i \ln(f) + (1-n_i) \ln(1-f)] \quad (747)$$

Written this way, $\Lambda(f; \{n_i\})$ define a rate of exponential growth for your winnings. But $\Lambda(f; \{n_i\})$ depends not only on your betting strategy, summarized by the fraction f that you put on heads, but also on the sequence of heads and tails that come up in the game, denoted by $\{n_i\}$. The key point is that, if we play *many* times, so we can think about the limit $N \rightarrow \infty$, this dependence on the details of the flips goes away.

We recall that, for any well behaved random variable, the average over N observations must approach the mean computed from the probability distribution as N becomes large. In the present case, if n_i is a binary variable that takes the value $n_i = 1$ with probability p and $n_i = 0$ with probability $1-p$, then as N becomes large we should have

$$\frac{1}{N} \sum_{i=1}^N n_i \rightarrow p, \quad (748)$$

and similarly

$$\frac{1}{N} \sum_{i=1}^N (1-n_i) \rightarrow 1-p. \quad (749)$$

We can use this to evaluate the long term growth of your winnings, simplifying the results of Eq (747):

$$\frac{1}{N} \ln G_{\text{total}}(N) \equiv \Lambda(f) = \ln 2 + \frac{1}{N} \sum_{i=1}^N [n_i \ln(f) + (1 - n_i) \ln(1 - f)] \quad (750)$$

$$\begin{aligned} &= \ln 2 + \left(\frac{1}{N} \sum_{i=1}^N n_i \right) \ln(f) + \left(\frac{1}{N} \sum_{i=1}^N (1 - n_i) \right) \ln(1 - f) \\ &\rightarrow \ln 2 + p \ln(f) + (1 - p) \ln(1 - f), \end{aligned} \quad (751)$$

where again p is the probability of heads. To maximize the growth rate $\Lambda(f)$, as usual we differentiate and set the result to zero:

$$\begin{aligned} \Lambda(f) &= \ln 2 + p \ln(f) + (1 - p) \ln(1 - f) \\ \frac{d\Lambda(f)}{df} &= p \frac{1}{f} + (1 - p)(-1) \frac{1}{1 - f}; \end{aligned} \quad (752)$$

$$\begin{aligned} \left. \frac{d\Lambda(f)}{df} \right|_{f=f_{\text{opt}}} &= 0 \\ \Rightarrow 0 &= p \frac{1}{f_{\text{opt}}} + (1 - p)(-1) \frac{1}{1 - f_{\text{opt}}} \end{aligned} \quad (753)$$

$$\frac{1 - p}{1 - f_{\text{opt}}} = \frac{p}{f_{\text{opt}}}, \quad (754)$$

or more simply $f_{\text{opt}} = p$. This is an interesting result: you maximize the rate at which your winnings will grow by “matching” the fraction of your resources that you bet on heads to the probability that the coin will come up heads, and similarly for tails.

Problem 144: Check that $f_{\text{opt}} = p$ is a maximum, and not a minimum, of $\Lambda(f)$.

Problem 145: If we bet only once, then in this simple game the maximum mean payoff is obtained by betting on the most likely outcome. On the other hand, as we play many times—more precisely, in the limit that we play infinitely many times—what we have seen is that a sort of matching strategy, or “proportional gambling” maximizes the growth rate. Explore the crossover between these limits. You might start with some simple simulations, and then see if you can make analytic progress, perhaps saying something about the leading $1/N$ corrections at large N . I am leaving this deliberately vague and open ended, hoping that you will play around.

Something even more interesting happens when we evaluate the optimal growth rate, that is $\Lambda_{\text{opt}} = \Lambda(f_{\text{opt}})$:

$$\Lambda_{\text{opt}} = \Lambda(f = p) \quad (755)$$

$$= \ln 2 + p \ln(p) + (1 - p) \ln(1 - p) \quad (756)$$

$$= \ln 2 - [-p \ln(p) - (1 - p) \ln(1 - p)]. \quad (757)$$

These terms should be starting to look familiar. The term $\ln 2$ is the entropy for a binary variable (heads/tails) if you don’t know anything about what to expect, and

hence the two alternatives are equally likely. In contrast, the term in brackets,

$$-p \ln(p) - (1 - p) \ln(1 - p),$$

is the entropy of a binary variable if you know that the two alternatives come up with probabilities p and $1 - p$. Thus the optimal growth rate is the difference in entropy between what might happen with an arbitrary coin and what you know will happen with this coin. In other words, *the maximum rate at which your winnings can grow in a simple gambling game is equal to the information that you have about the outcome of a single coin flip.*

This connection between information theory and gambling was discovered in the 1950s by Kelly, who was searching for some interpretation of Shannon’s work that didn’t refer to the process of communication. Obviously what we have worked out here is a very simple and special case, and we need to do much more in order to claim that the connection is general. But before launching into this let me emphasize something about Kelly’s result. At some intuitive level, we can all agree that if we know more about the outcome of the coin flip (or the horse race, or the stock market, or ...) then we should be able to make more money. In a very general context, Shannon proved that “know more” should be quantified by various entropy-like quantities, but it’s not obvious that the knowledge measured by Shannon’s bits is actually the useful knowledge when it comes time to make a bet. Even if bits are the right measure, the connection between information and the growth of winnings could have been much more vague; you could imagine, for example, that the growth rate is bounded by some function of the information, and that this bound might or might not be realizable with feasible strategies. In contrast to these pessimistic alternatives, Kelly showed that the maximum growth rate *is* the information, and his proof is constructive so we actually know how to achieve this maximum. This really is quite astonishing.

Let’s try to generalize what we have done. Suppose that on each trial i , there are many possible outcomes, $\mu = 1, 2, \dots, K$; we’ll write $n_i^{(\mu)} = 1$ if on the i^{th} trial the outcome is μ , and $n_i^{(\mu)} = 0$ otherwise. Further, let’s say that you bet a fraction of your assets f_μ on each of the possible outcomes μ , and if μ actually happens then each dollar bet on this outcome becomes g_μ dollars;

all money bet on things that don't happen is lost. If you need an example of this sort of game, think of a horse race in which you get something back only if you pick the winner. We'll assume that the different outcomes occur with probability p_μ , but we won't assume anything about the relationship between these odds and the payoffs g_μ .

Having defined all the factors, the analog of Eq (743), is

$$G_{\text{total}}(N) = \prod_{i=1}^N \left[\sum_{\mu=1}^K f_\mu g_\mu n_i^{(\mu)} \right]. \quad (758)$$

Now we can follow the same steps as before:

$$\ln G_{\text{total}}(N) = \sum_{i=1}^N \ln \left[\sum_{\mu=1}^K f_\mu g_\mu n_i^{(\mu)} \right] \quad (759)$$

$$= \sum_{i=1}^N \sum_{\mu=1}^K n_i^{(\mu)} \ln(f_\mu g_\mu) \quad (760)$$

$$\frac{1}{N} \ln G_{\text{total}}(N) = \sum_{\mu=1}^K \left[\frac{1}{N} \sum_{i=1}^N n_i^{(\mu)} \right] \ln(f_\mu g_\mu) \quad (761)$$

$$\rightarrow \Lambda(\{f_\mu\}) = \sum_{\mu=1}^K p_\mu \ln(f_\mu g_\mu). \quad (762)$$

We want to maximize the growth rate Λ , subject to the normalization condition that the fractions of our assets placed on all the options add up ($\sum_{\mu} f_\mu = 1$), so we introduce a Lagrange multiplier α and find the maximum of the function

$$\tilde{\Lambda}(\{f_\mu\}) = \sum_{\mu=1}^K p_\mu \ln(f_\mu g_\mu) - \alpha \left[\sum_{\mu=1}^K f_\mu - 1 \right]. \quad (763)$$

The equations for the maximum are, as usual,

$$\left. \frac{\partial \tilde{\Lambda}(\{f_\mu\})}{\partial f_\mu} \right|_{\{f_\mu\}=\{f_\mu^{\text{opt}}\}} = 0 \quad (764)$$

$$\Rightarrow 0 = \frac{p_\mu}{f_\mu^{\text{opt}}} - \alpha, \quad (765)$$

$$f_\mu^{\text{opt}} = \frac{p_\mu}{\alpha}; \quad (766)$$

since $\sum_{\mu} f_\mu = \sum_{\mu} p_\mu = 1$, we must have $\alpha = 1$, so that

$$f_\mu^{\text{opt}} = p_\mu. \quad (767)$$

Substituting, we find the maximum growth rate

$$\Lambda_{\text{opt}} = \sum_{\mu=1}^K p_\mu \ln(p_\mu g_\mu). \quad (768)$$

The first interesting thing is that we recover from the simpler heads/tails problem the idea of proportional gambling [Eq (767)]: you maximize the rate at which

your winnings will grow by “matching” the fraction of your resources that you bet on each horse in the race to the probability that this horse will win. Strangely, this is independent of the rewards or gains as expressed in the parameters $\{g_\mu\}$.

[At some point should make a connection between proportional gambling and “matching” behavior .. is this understood?]

The second point is that we can see what it means for the odds to be truly fair. If our opponent in this game (the track operator) sets the returns in inverse proportion to the probability that each horse wins, $g_\mu = 1/p_\mu$, then the maximum growth rate of our winnings, Λ_{opt} , is exactly zero.

This notion of fairness leads us to an information theoretic interpretation of Λ_{opt} . Notice that we have done our calculation on the assumption that we have perfect knowledge of the distribution $\{p_\mu\}$. Perhaps the track operators have less knowledge, and so they set the odds *as if* the distribution were something else, which we can call $\{q_\mu\}$. More generally, we can define

$$q_\mu = \frac{1}{Z} \frac{1}{g_\mu}, \quad (769)$$

with Z chosen so that $\sum_{\mu} q_\mu = 1$. If $Z = 1$, then the payoffs $\{g_\mu\}$ are fair in the distribution $\{q_\mu\}$, while if $Z < 1$ the track operators are keeping something for themselves (as they are wont to do). Then we can see that

$$\Lambda_{\text{opt}} = -\ln Z + \sum_{\mu=1}^K p_\mu \ln \left(\frac{p_\mu}{q_\mu} \right). \quad (770)$$

You should recognize the second term as the Kullback-Leibler divergence between the probability distributions $\mathbf{p} \equiv \{p_\mu\}$ and $\mathbf{q} \equiv \{q_\mu\}$, from Eq (656).

$$D_{\text{KL}}(\mathbf{p}||\mathbf{q}) \equiv \sum_{\mu=1}^K p_\mu \ln \left(\frac{p_\mu}{q_\mu} \right). \quad (771)$$

We recall that the KL divergence measures the cost of coding signals with the wrong distribution. Equation (770) shows us that better knowledge of the probability distribution doesn't just allow us to make shorter codes. The amount by which we can compress the data describing the sequence of winners in the horse race is exactly the amount by which our winnings can grow. More precisely, if we can build a shorter code than the one built implicitly by the track operators, then we will gain exactly in proportion to this shortening. Thus, in this context, we literally get paid for constructing more efficient representations of the data (!).

We have connected the growth rate of winnings to the efficiency with we can represent data, but this isn't quite as compelling as a direct connection to how much information we have about the outcome of the game, which

is where we started in the case of coin flips; let's see if we can do better. Imagine that, on each trial i , we have access to some signal x_i that tells us something about the likely outcome. More precisely, when we observe x_i , the probability that the outcome will be μ on trial i is not p_μ but rather some conditional probability $p(\mu|x_i)$; if the signals x are themselves chosen from some distribution $P(x)$, then for consistency we must have

$$p_\mu = \int dx P(x) p(\mu|x). \quad (772)$$

To use the extra information provided by the signal x , you will adjust your strategy to bet a fraction $f_\mu(x_i)$ on the outcome μ given that you have 'heard' x_i . How does the extra information provided by x improve your winnings?

To compute the growth of winnings in the presence of extra information, we proceed along the same lines as before, to find the analog of Eq (762):

$$\Lambda[\{f_\mu(x)\}] = \int dx P(x) \sum_{\mu=1}^K p(\mu|x) \ln[f_\mu(x)g_\mu]. \quad (773)$$

Now we need to maximize this, choosing strategies that are defined by the *functions* $f_\mu(x)$, where for each x we have the constraint that $\sum_\mu f_\mu(x) = 1$. Once again the solution to this optimization problem is proportional gambling, but now the proportions are conditioned on your knowledge, so that the analog of Eq (767) becomes

$$f_\mu^{\text{opt}}(x) = p(\mu|x). \quad (774)$$

This determines the optimal growth rate,

$$\Lambda_{\text{opt}} = \int dx P(x) \sum_{\mu=1}^K p(\mu|x) \ln[p(\mu|x)g_\mu]. \quad (775)$$

Problem 146: Fill in the steps leading to the derivation of $\Lambda[\{f_\mu(x)\}]$ in Eq (773) and the consequences of optimizing this functional, Eq's (774) and (775).

The important result is the gain in growth rate that is possible by virtue of having access to the signal x , that

is the difference between Λ_{opt} in Eq (775) and Eq (768):

$$\Delta\Lambda_{\text{opt}} = \int dx P(x) \sum_{\mu=1}^K p(\mu|x) \ln[p(\mu|x)g_\mu] - \sum_{\mu=1}^K p_\mu \ln[p_\mu g_\mu] \quad (776)$$

$$= \int dx P(x) \sum_{\mu=1}^K p(\mu|x) \ln[p(\mu|x)g_\mu] - \int dx P(x) \sum_{\mu=1}^K p(\mu|x) \ln[p_\mu g_\mu] \quad (777)$$

$$= \int dx P(x) \sum_{\mu=1}^K p(\mu|x) \ln \left[\frac{p(\mu|x)}{p_\mu} \right]. \quad (778)$$

We see that the details of the payoffs g_μ drop out, and that *the gain in growth rate is exactly the mutual information between the signal x and the outcomes μ .*

Once again information translates directly into the (increased) rate at which capital can grow. Thus, the abstract measure of information has a clear impact on very down to earth measures of performance in a real world task. But, beyond metaphor,⁸³ what does this have to do with life?

The most direct connection between life and gambling is through the phenomenon of persistence. Many bacteria have two distinct lifestyles. In one (for example), they grow quickly in most environments, but are very susceptible to being killed by antibiotics. In the other, they grow very slowly, but survive the antibiotics. This is almost exactly the horse race—if the bacterium bets correctly, it grows, but if it bets incorrectly it dies (or grows at rates far below what is possible). Absent any direct measurements on the environment, a population of genetically identical bacteria will maximize its growth rate by a form of proportional gambling, so that even in a healthy person, not taking antibiotics, we should see that some of the resident bacteria persist in a state of slow growth and (eventual) antibiotic resistance;⁸⁴ the fraction of bacteria in this states reflects the population's estimate of the probability that they will encounter the hostile environment of antibiotics [do we know anything about whether bacteria are doing this correctly?]. We also see that gaining information about the environment opens the possibility of faster growth, in precise proportion to the information gained.

⁸³ Life is a gamble, etc..

⁸⁴ Here "resistance" is used colloquially. Technically, antibiotic resistance refers to a trait which is encoded genetically, and hence inheritable, rather than a lifestyle choice. The (choosable) state in which bacteria grow slowly but are not killed by antibiotics is called "persistent."

In a world of two alternatives, there is not much information to gain. There are examples of bacteria that choose among a wider variety of lifestyles, and these phenomena (including the simple example of two alternatives) are called ‘phenotypic switching.’ In the approximation that for each environment there is only one phenotype which grows, phenotypic switching is exactly the horse racing problem.

Problem 147: Something based on phenotypic switching .. look through Kussell et al for ideas.

The example of phenotypic switching makes a nice map back to the early work about gambling, but is perhaps still a bit too simple. Let’s try to be more general. Imagine a bacterium that lives in an environment in which the concentrations of nutrients are fluctuating (slowly, so we don’t have to worry about dynamics). In order to make use of the currently available nutrients, the bacterium must express the relevant enzymes involved in metabolism. Let’s simplify and assume that there is one nutrient or substrate at concentration s and one relevant gene at expression level g . The bacterium will then grow at some rate $r(s, g)$ that depends both on the state of the world (s) and on its internal state (g).

The growth rate of the bacterium is a compromise between two effects. On the one hand, growth requires metabolism of the available nutrient, and so growth should be faster if there is either more nutrient or more enzyme. On the other hand, making the enzyme itself takes resources, and this should slow the growth; in the limit of small nutrient concentrations, this cost can become dominant, and growth would stop if the cell tried to make too much enzyme. This scenario is shown schematically in Fig 140.

Problem 148: A simple fitness landscape. The schematic in Fig 140 is based on a simple model. Suppose that growth is precisely proportional to the rate at which the enzyme degrades the substrate. In a Michaelis–Menten kinetic scheme for the enzyme [pointer to earlier discussion of MM kinetics], this means that the rate of degradation (in molecules per second) will be

$$V = V_{\max} g \frac{s_{\text{free}}}{K + s_{\text{free}}}, \quad (779)$$

where g is the number of copies of the enzyme molecule, V_{\max} is the maximum rate at which the enzyme can run, s_{free} is the concentration of the substrate free in solution, and K is the ‘Michaelis constant’ that sets the scale for half-saturation of the enzyme. The total substrate concentration is the sum of that free in solution and bound to the enzyme,

$$s = s_{\text{free}} + \frac{1}{\Omega} g \frac{s_{\text{free}}}{K + s_{\text{free}}}, \quad (780)$$




FIG. 140 A schematic of bacterial growth rate as a function of available substrate concentration and enzyme expression level. The growth rate is a compromise between metabolizing the substrate and the cost of making the enzyme. The thin white line [redraw!] traces the optimal setting of expression level as a function of substrate availability.

where Ω is the cell volume. If the growth rate is proportional to the metabolic rate, less a correction for the cost of making the enzymes, we should have

$$r(s, g) = \alpha g \frac{s_{\text{free}}}{K + s_{\text{free}}} - \beta g. \quad (781)$$

Solve for s_{free} to rewrite $r(s, g)$ explicitly in terms of s . Then show that by proper choice of units, there is only one arbitrary parameter. What is the meaning of this remaining parameter? Make some reasonable choices, and plot your own version of Fig 140.

Imagine a bacterium whose life is governed by Fig 140. As the available substrate concentration fluctuates, one possibility is that all bacteria carefully adjust their enzyme expression levels to achieve optimal growth rate under each condition. An extreme alternative is that different bacteria in the population choose their expression levels at random out of some distribution, and hope that some of them by chance have made good choices, much as in the proportional gambling scenario. In the first case, the expression level carries an enormous amount of information about the concentration of available substrate—indeed, if we imagine that the optimum is traced perfectly, then knowing the expression level would tell us the exact substrate concentration, and this represents an infinite amount of information (!). In contrast, the gambling strategy involves no correlation of the internal and external states, and hence no information is conveyed.

Evidently, the average growth rate across an ensemble of environments will be larger if the bacteria can adjust their expression levels perfectly, but maybe this is so obvious as not to be interesting. We know that there is some average growth rate which can be achieved with no information about the outside world, and that an infinite amount of information would allow the population to grow faster. What happens in between?

The mutual information between the internal state g and the external world s can be written as

$$I(g; s) = \int ds P(s) \int dg P(g|s) \log_2 \left[\frac{P(g|s)}{P(g)} \right]. \quad (782)$$

We can make $I(g; s)$ as small as we like by letting $P(g|s)$ approach $P(g)$. But suppose that we want to maintain some average growth rate in the ensemble of environments defined by $P(s)$. This average growth rate is

$$\langle r \rangle = \int ds P(s) \int dg P(g|s) r(s, g). \quad (783)$$

Now it seems clear that not all conditional distributions $P(g|s)$ are consistent with a given $\langle r \rangle$. What we would like to show is that there is a minimum value of $I(g; s)$ consistent with $\langle r \rangle$.

The problem we have is a constrained minimization, so as usual we introduce a Lagrange multiplier and minimize

$$\mathcal{F}[P(g|s)] \equiv I(g; s) - \lambda \langle r \rangle - \int ds \mu(s) \int dg P(g|s), \quad (784)$$

where the second set of Lagrange multipliers $\mu(s)$ enforces normalization of the distributions $P(g|s)$ at each value of s . Finding the minimum in this case is straightforward. The key step is to evaluate the derivative of the information with respect to the conditional distribution:

$$\frac{\delta I(g; s)}{\delta P(g|s)} = \frac{\delta}{\delta P(g|s)} \int ds P(s) \int dg P(g|s) \log_2 \left[\frac{P(g|s)}{P(g)} \right] \quad (785)$$

$$= P(s) \log_2 \left[\frac{P(g|s)}{P(g)} \right] + \frac{1}{\ln 2} P(s) P(g|s) \cdot \frac{1}{P(g|s)} - \frac{1}{\ln 2} \int ds' P(s') P(g|s') \frac{1}{P(g)} \frac{\delta P(g)}{\delta P(g|s)} \quad (786)$$

$$= P(s) \log_2 \left[\frac{P(g|s)}{P(g)} \right] + \frac{1}{\ln 2} P(s) - \frac{1}{\ln 2} P(g) \frac{1}{P(g)} P(s) \quad (787)$$

$$= P(s) \log_2 \left[\frac{P(g|s)}{P(g)} \right], \quad (788)$$

which is nice because all the messy bits cancel out. Now we can solve our full problem:

$$0 = \frac{\delta \mathcal{F}[P(g|s)]}{\delta P(g|s)} \quad (789)$$

$$= \frac{\delta}{\delta P(g|s)} \left[I(g; s) - \lambda \int ds P(s) \int dg P(g|s) r(s, g) - \int ds \mu(s) \int dg P(g|s) \right] \quad (790)$$

$$= P(s) \log_2 \left[\frac{P(g|s)}{P(g)} \right] - \lambda P(s) r(s, g) - \mu(s) \quad (791)$$

$$\log_2 \left[\frac{P(g|s)}{P(g)} \right] = \lambda r(s, g) + \frac{\mu(s)}{P(s)} \quad (792)$$

$$P(g|s) = \frac{1}{Z(s)} P(g) \exp[\beta r(s, g)], \quad (793)$$

where $\beta = \lambda \ln 2$, and $Z(s) = \exp[\ln 2 \mu(s)/P(s)]$ is a normalization constant,

$$Z(s) = \int dg P(g) \exp[\beta r(s, g)], \quad (794)$$

and of course we must obey

$$P(g) = \int ds P(s) P(g|s). \quad (795)$$

Notice that our solution for $P(g|s)$ is (roughly) a Boltzmann distribution, where $-r(s, g)$ plays the role of the energy and β is the inverse temperature. As expected




FIG. 141 Mean growth rate as a function of the mutual information between expression levels and substrate availability for the system in Fig 140. We assume that the (log) substrate is chosen from a distribution that is uniform over the 16-fold range shown in Fig 140, and then we solve for the optimal $P(g|s)$ using Eq's (793–794).

from this analogy, we can write the information and average growth rate as derivatives,

$$I(g; s) = \lambda \langle r \rangle - \int ds P(s) \log_2 Z(s), \quad (796)$$

$$\langle r \rangle = \int ds P(s) \frac{d \ln Z(s)}{d\beta}. \quad (797)$$

The Boltzmann form of the optimal solution in Eq (793) helps our intuition. At small β , the distribution $P(g|s)$ is almost the same as $P(g)$, so that very little information is conveyed between internal and external states. In contrast, as λ becomes large, the distribution $P(g|s)$ becomes sharply peaked around the value expression level $g_{\text{opt}}(s)$ that maximizes the growth rate. Varying β should trace out a curve of mean growth rate vs. information, and this is shown in Fig 141. We see from the derivation that this curve represents the maximum mean growth rate achievable given a certain amount of mutual information, or alternatively the minimum amount of information required to achieve a certain mean growth rate, $I_{\text{min}}(\langle r \rangle)$.

Problem 149: Asymptotics of growth rate vs information. The precise form of the relationship between the mean growth rate and the minimum information depends, of course, on details of the function $r(s, g)$. Show that the behavior at large values of the minimum information is more nearly universal. To do

this, develop an asymptotic expansion at large values of λ ,

$$\begin{aligned} P(g|s) &= \frac{1}{Z(s)} P(g) \exp \left[\tilde{\lambda} r(s, g) \right] \\ &\approx \frac{1}{Z(s)} P(g) \exp \left[\tilde{\lambda} r(s, g_{\text{opt}}(s)) + \frac{\tilde{\lambda}}{2} A (g - g_{\text{opt}}(s))^2 \right], \end{aligned} \quad (798)$$

$$A = \left. \frac{\partial^2 r(s, g)}{\partial g^2} \right|_{g=g_{\text{opt}}(s)} \quad (799)$$

and use this expansion to evaluate $Z(s)$, from which you can calculate $I_{\text{min}}(\langle r \rangle)$. Can you generalize your discussion to the case where there are many substrates and many genes to control?

It is important to take seriously the scales in Fig 141. It could have been that the full growth advantage derived from controlling expression levels was achievable with only a small fraction of a bit, or conversely that it required many tens of bits. In fact, for this simple problem the answer is that cells can make use of more than one bit, but not too much more. This means that (near-)optimal growth requires more than just turning a gene on and off, and presumably this is even more clear if we think about more realistic situations where there are multiple substrates and multiple genes. As we will see in the next section, the noise levels measured for the control of gene expression set a limit of $\sim 1 - 3$ bits to the information that can be transmitted through these control elements. Thus, the amount of information that cells need in order to optimize their growth in varying environments is plausibly close to the maximum they can transmit, and this limit in turn is set by the number of molecules that the cell is devoting to these tasks.

Just to be clear, it's useful to think about the alternatives. If information is cheap, so that it is easy for cells to transmit many bits, then evolution selects for mechanisms that drive the system upward in the information/fitness plane of Fig 141. But if information itself is hard to come by, evolutionary pressure (which really only acts to increase growth rates) must necessarily drive cells outward along the information axis.

Sometimes the fact that organisms have to be flexible and survive in a fluctuating environment is offered as a qualitative argument against the possibility of optimization. Indeed, if the environment fluctuates, it may not be advantageous for organisms to drive toward “perfect” performance under any one set of conditions. But the argument we have given here shows that strategies for dealing with varied environments are themselves subject to optimization, making the most of a limited amount of information and eventually being pushed by selection to gather more bits.

In the problem of horse races, or phenotypic switching, information translated directly into a growth rate. Here

we see that, more generally, there is a minimum amount of information needed to achieve a given average growth rate. In both of these cases, information is necessary and permissive, but not sufficient. Thus organisms *can* grow faster if they gather and represent more information, but this is not guaranteed—they might make poor use of the information, and fail to reach the bound on their growth rate. We have focused here on achieving a certain average growth rate, but it should be clear that the whole discussion can be transposed to other domains. For example, if I ask you to point at a target that can appear at random in your visual field, and reward you in proportion to how close you come to the exact position of the target, then in order to collect a certain level of average reward your brain must represent some minimum amount of information about the target location. Quite generally, we can imagine plotting some “biological” measure of performance—probability of catching a mate, nutritional value extracted from picking fruit, growth rate, happiness, ... —versus the amount of information that the organism has about the relevant variables. This “information/fitness” plane will be divided by a curve which separates the possible from the impossible, since without a certain minimum level of information, higher fitness is impossible.

Problem 150: Information and motor control. Give a simple example, maybe from smooth pursuit?

In the information theory literature, the sort of bounds we are computing here go by the name of “rate–distortion” curves. For example, if we measure image quality by some complicated perceptual metric, then to have images of a certain quality, on average, we will need to transmit a minimum number of bits. In this spirit, we can think about more complicated situations, such as organisms foraging or acting in response to sensory stimuli and collecting rewards. Although one is not rewarded specifically for bits, the message of rate–distortion theory is that to collect rewards at some desired rate will always require a minimum number of bits of information.

In constructing a rate distortion curve, we implicitly define some bits as being more relevant than others. Thus if I need to match my state to that of the environment, presumably some environmental variables need to be tracked more accurately than others; since the rate distortion curves gives the minimum number of bits, I need to get this right and put the precision (extra bits) in the right place. This is important, because it means that we have a framework for assigning value to bits. To be concrete, in Fig 794 it is possible to imagine an infinite

variety of mechanisms that gather the same number of bits but fail to achieve the maximum mean growth rate, either because they use the bits incorrectly or because they have gathered the wrong bits. Bits in and of themselves are not guaranteed to be useful, but to do useful things there is a minimum number of bits that we need.

An interesting if unfinished connection of rate–distortion theory to biological systems is the case of protein structure. If I want to describe protein structures with high precision, I need to tell you where every atom is located. But if sequence determines structure, then to some accuracy I just need to tell you the amino acid sequence, which is at most $\log_2(20)$ bits per amino acid, and many fewer per atom. In fact, as we have discussed in Section III.A, many different sequences generate essentially the same structure, so there must be an even shorter description. Thus, if we imagine taking the ensemble of real protein structures, there must be a description in very few bits that nonetheless generates rather small errors in predicting the positions of the atoms. Finding the optimally compact description (i.e., along the true rate–distortion curve) would be a huge help in understanding protein folding, because the joint table of sequences and (compactly described) structures would be much smaller. There is even an intuition that there must be such a compact representation with high accuracy in order to make folding rapid, essentially because the number of states needed for an accurate description should be connected to the number of states that the protein must “search” through as it folds. I am not sure how to make this rigorous, but it’s interesting.

Problem 151: Clustering structures. Give an example of constructing rate–distortion curves via clustering ... maybe something plausibly connected to molecular structures?

We can search for compact descriptions of protein structure by approximating the local path of the α –carbon backbone as moves on a discrete lattice, making the lattice progressively more complex. We can do better by moving off the lattice to cluster the natural dihedral angles describing the path from one amino acid to the next **Be sure we talk about ϕ, ψ description of proteins before this, and point back**]; results are shown in Fig 142. Indeed, by the time we have assigned 10 or 20 states per amino acid, we can reconstruct structures with 1 – 2 Å rms accuracy.

Another very specific connection between biology and bits is in the case of embryonic development. In the simplest model of morphogen gradients, each independently “reads out” the local concentration of the morphogen(s),




FIG. 142 Rate–distortion curve (or its moral equivalent) for protein structures, from Park and Levitt (1995). The path of the α -carbon backbone is approximated by a discrete set of local ‘moves’ along the chain, which is roughly equivalent to forcing the structure to live on a lattice. Diamonds correspond to lattices with different structures (e.g., 3 possible moves on a tetrahedral lattice); + and W correspond to discrete approximations obtained by clustering known structures based on the Ramachandran angles at each site. Plotted on the y-axis is the root mean square error in the positions of all the α -carbons along the chain. Inset shows two examples of protein structures, compared with their discrete approximations.

and makes decisions—most importantly, about the regulation of gene expression—based on this local measurement, as in Fig 143. In this model, the only thing that a cell knows about its position in the embryo is the morphogen concentration, and so the information that cells have about position can be no larger than the information that they extract about this concentration. In effect there is a communication channel from the morphogen to the expression levels of the genes which defines the blueprint for development, and the information that can be transmitted along this channel sets a bound on the complexity and reliability of the blueprint. As an example, if we have N rows of cell along one axis of the embryo, and each row reliably adopts a distinct fate that we can ‘see’ by looking at the expression levels of a handful of genes, then (again, in the simplest model) there must be $\log_2 N$ bits of information transmitted through the regulatory network that takes the morphogens as input and gives the gene expression levels as output. As in the discussion of growth rates, this becomes interesting because, as we shall see, the information capacity of gene regulatory elements is quite limited. Rough estimates of the relevant quantities in the *Drosophila* embryo suggest that the embryo might indeed be forming patterns near the limits set by the information capacity of gene regulation.

What happens if things are more complicated than in Fig 143? In particular, we know about plenty of systems which form patterns spontaneously, without any analog of the “maternal” signal to break the translational symmetry. It is important to realize that while patterns can form spontaneously, information can’t really be created, only transmitted. In a crystal, for example, once we know that one atom is in a particular position we can predict the position of other atoms, but this is only because of the bonds that connect the atoms. Because all the atoms undergo Brownian motion, the transmission of information is not perfect, and knowledge of one atomic position provides only a limited number of bits about the position of another atom; this limit on information transmission becomes tighter as the temperature—and hence the noise level in the “communication channel” which connects the distant atoms—becomes larger, until the crystal melts and there is no information transmitted over long distances.

Problem 152: Transmitting positional information in a crystal. Take the students through an explicit calculation of the mutual information between positions of atoms in a harmonic solid.

In non-equilibrium systems, such as the Rayleigh–Bernard convection cell shown in Fig 144, we see spatial patterns in which some local variable such as the temperature, fluid density or velocity at one position predicts




FIG. 143 Information flow in a “feed–forward” model of genetic control in the early embryo. The concentration C of the primary morphogen depends on position x , and each cell responds independently by modulating the expression level G of some target gene (or genes). In this simple view, information about position only reaches the gene expression level through the intermediary of the primary morphogen concentration, and hence we have $I(x; G) \leq I(C; G)$.




FIG. 144 This looks like a perfect crystal of beads, but it actually is a small (~ 10 cm diameter) container filled with carbon dioxide at high pressure, and heated from below. The image is formed by passing light through the gas, sometimes called a ‘shadowgraph.’ The temperatures at the top and bottom of the container are held very constant (to within a few thousandths of a degree) so that the patterns will not be disrupted by variations in conditions; similarly, the top and bottom of the container are extremely flat (smooth to within the wavelength of light), and the whole system is held horizontal with high precision so that the direction of gravity is aligned with the axis of symmetry through the center of the circle. From E Bodenschatz et al (1991).

the value of the corresponding variable at another position. If we call this local variable $\phi(x)$, then if we imagine a large ensemble of snapshots like the one in Fig 144, we can build up the distribution functional $P[\phi(\vec{x})]$. The statement that we have a periodic pattern, for example, is the statement that if we look at two points separated by an appropriately chosen vector \vec{d} , then $\phi(\vec{x}) \approx \phi(\vec{x} + \vec{d})$. But if we point to the first point \vec{x} at random, we can get a broad range of values for $\phi_1 \equiv \phi(\vec{x})$, drawn from a distribution $P_1(\phi_1)$. Similarly, if we are choosing \vec{x} at random then $\phi_2 \equiv \phi(\vec{x} + \vec{d})$ is also broadly distributed; in fact, it must come from the same distribution as ϕ_1 . But once we know ϕ_1 , if there is a periodic pattern then the distribution $P(\phi_2|\phi_1)$ must be sharply peaked around $\phi_1 = \phi_2$, and hence very different from the “prior” distribution of ϕ_2 . But this is exactly the condition for there to be mutual information between ϕ_1 and ϕ_2 . Thus, the existence of a spatial pattern is equivalent to the presence of mutual information between the local variables at distant points. Where does this information come from? As with the bonds connecting the atoms in the crystal, it must be transmitted through the dynamics of the system, which connect points only to their immediate neighbors.

In a strict interpretation of the concept of positional information in embryo, we actually require more than mutual information between local variables at distant

points. We require that the value of some local variable(s), typically the expression levels of several genes, tell us about the location of the point where we have observed them. In this way, cells would “know” their position in the embryo by virtue of their expression levels, and these signals could drive further processes in a way that is appropriate to the cell’s location—not just relative to other cells, but in absolute terms.⁸⁵ If we call the local variables $\{g_i\}$, for gene expression levels, then the positional information is $I(\vec{x}; \{g_i\})$. But the local variables at point x are controlled by a set of inputs which may include external, maternally supplied morphogens, the expression levels $\{g_i\}$ in neighboring cells, and perhaps other variables as well. We can always write the distribution of expression levels at one point in terms of this inputs,

$$P(\{g_i\}|x) = \int d(\text{inputs}) P(\{g_i\}|\text{inputs})P(\text{inputs}|x). \quad (800)$$

Noise in the control of gene expression corresponds to the fact that the distribution $P(\{g_i\}|\text{inputs})$ is not infinitely narrow. Now because, at any one point, information flows from x to the inputs to the $\{g_i\}$, we must have $I(x; \{g_i\}) \leq I(\text{inputs}; \{g_i\})$, and this is true no matter how complicated the inputs might be. More importantly, as hinted at in the analysis of the first synapse in fly vision (Fig 139), any input/output device has a maximum amount of information it can transmit that is determined by its noise level. Thus, if we think of all the whole network of interactions that result in the regulation of the gene expression levels $\{g_i\}$, the noise in this network determines a maximum value for $I(\text{inputs}; \{g_i\})$, and this sets a limit to the amount of positional information that cells in the embryo can acquire and encode with these genes.

Problem 153: The data processing inequality. What we need in the previous paragraph is a special case of a more general inequality .. derive it.

To summarize, the reliability and complexity of the patterns that can form during embryonic development

⁸⁵ This is certainly what “positional information” means in the usual descriptions of the concept; see the discussion of the information carried by Hunchback expression levels in the fly embryo, surrounding Fig 136. There are almost no measurements of this information, in bits, so it remains possible that real cells know much more about their relative position than about their absolute position. This wouldn’t change the spirit of what I am saying here, but the details would matter. This is one of many open questions about information flow in the embryo.

are limited by the amount of positional information that cells can acquire and represent. This information in turn is limited by the “capacity” of the genetic or biochemical networks whose outputs encode the positional information. Therefore, if real networks operate in a regime where this capacity is small, the complexity of body plans will be limited by the ability of the organism to squeeze as much information as possible out of these systems.

Most of the examples we have considered thus far have the feature that the information is “about” something that has obvious relevance for the organism. Can we find some more general way at arriving at such notions of relevant? It is useful to have in mind an organism collecting a stream of data, whether the organism is like us, with eyes and ear, or like a bacterium, sensing the concentrations of various molecules in its external and internal environment. Of all these data, the only part we can use to guide our actions (and eventually collect rewards, reproduce, etc.) is the part that has predictive power, since by the time we act we are already in the future. Thus we can ask how to squeeze, out of all the bits we collect, only those bits which are relevant for prediction.




FIG. 145 A schematic of the prediction problem. We observe a time series, and at some moment (now) we look back at a segment of the recent past with duration T , X_{past} . From this, we try to infer something about what will evolve in the future.

[I am worried that this goes a little quickly.] More concretely, as in Fig 145, if we observe a time series through a window of duration T (that is, for times $-T < t \leq 0$), then to represent the data X_{past} we have collected requires $S(T)$ bits, where S is the entropy, but the infor-

mation that these data provide about the future X_{future} (i.e., at times $t > 0$) is given by some $I(X_{\text{past}}; X_{\text{future}}) \equiv I_{\text{pred}}(T) \ll S(T)$. In particular, while for large T the entropy $S(T)$ is expected to become extensive, the predictive information $I_{\text{pred}}(T)$ always is subextensive. Thus we expect that the data X_{past} can be compressed significantly into some internal representation X_{int} without losing too much of the relevant information about X_{future} . Formally, we can construct the optimal version of this mapping by solving

$$\max_{X_{\text{past}} \rightarrow X_{\text{int}}} [I(X_{\text{int}}; X_{\text{future}}) - \lambda I(X_{\text{int}}; X_{\text{past}})], \quad (801)$$

where $X_{\text{past}} \rightarrow X_{\text{int}}$ is the rule for creating the internal representation and λ is a Lagrange multiplier. This sort of problem has been dubbed an ‘information bottleneck,’ because we try to preserve the relevant information while squeezing the input data through a narrow channel.

Problem 154: Predictive information is subextensive.

If we observe a stationary stochastic process, $x(t)$, on the interval $t_1 < t \leq t_1 + T$, the entropy of the distribution $P[x(t)]$ depends only on T , not t_1 ; let’s call this entropy $S(T)$.

(a.) Use your intuition from statistical mechanics to explain why we expect $S(T)$ to grow extensively, that is $S(T) \propto T$ at large T . More formally, show that at large T

$$S(T) \rightarrow ST + S_1(T), \quad (802)$$

where

$$\lim_{T \rightarrow \infty} \frac{S_1(T)}{T} = 0. \quad (803)$$

Thus, although $S_1(T)$ can grow with T , it must grow more slowly than T itself—it is “subextensive.”

(b.) Consider the case where time is discrete, and x is Markovian, so that $x(t+1)$ depends on $x(t)$, but no earlier history. Show that, in this case, $S_1(T)$ is just a constant.

(c.) Consider the case where $X_{\text{past}} \equiv x(-T < t \leq 0)$ and $X_{\text{future}} \equiv x(0 < t < T')$. Show how the predictive information $I_{\text{pred}}(T, T') \equiv I(X_{\text{past}}; X_{\text{future}})$ is related to the function $S(T)$; you should be able to do this in general, without the Markov assumption. Show further that there a finite limit as the duration of the future becomes infinite, and that this limit $I_{\text{pred}}(T)$ is subextensive.

In general, we should consider the mapping $X_{\text{past}} \rightarrow X_{\text{int}}$ to be probabilistic, so we can describe it by some conditional distribution $P(X_{\text{int}}|X_{\text{past}})$. Then the quantity we are trying to maximize becomes

$$\begin{aligned}
-\mathcal{F} = & \sum_{X_{\text{int}}, X_{\text{past}}} P(X_{\text{int}}|X_{\text{past}})P(X_{\text{past}}) \log_2 \left[\frac{P(X_{\text{int}}|X_{\text{past}})}{P(X_{\text{int}})} \right] \\
& - \lambda \sum_{X_{\text{int}}, X_{\text{future}}} P(X_{\text{int}}|X_{\text{future}})P(X_{\text{future}}) \log_2 \left[\frac{P(X_{\text{int}}|X_{\text{future}})}{P(X_{\text{int}})} \right]. \tag{804}
\end{aligned}$$

This is written as if our choice of representation X_{int} depends directly on the future, but of course this isn't true; any correlation between what we write down and what happens in the future is inherited from the data that we collected in the past,

$$P(X_{\text{int}}|X_{\text{future}}) = \sum_{X_{\text{past}}} P(X_{\text{int}}|X_{\text{past}})P(X_{\text{past}}|X_{\text{future}}). \tag{805}$$

In addition, we have

$$P(X_{\text{int}}) = \sum_{X_{\text{past}}} P(X_{\text{int}}|X_{\text{past}})P(X_{\text{past}}). \tag{806}$$

As usual, we have to take the derivative of \mathcal{F} with respect to the distribution $P(X_{\text{int}}|X_{\text{past}})$, being careful to add a Lagrange multiplier $\mu(X_{\text{past}})$ that fixes the normalization for each value of X_{past} , and then we set the derivative to zero to find an extremum. Since the optimization of \mathcal{F} is independent of multiplicative factors, we can make things simpler by taking natural logs instead of logs base 2. Then the algebra is as follows:

$$0 = \frac{\delta}{\delta P(X_{\text{int}}|X_{\text{past}})} \left[-\mathcal{F} - \sum_{X_{\text{past}}} \mu(X_{\text{past}}) \sum_{X_{\text{int}}} P(X_{\text{int}}|X_{\text{past}}) \right] \tag{807}$$

$$= P(X_{\text{past}}) \ln \left[\frac{P(X_{\text{int}}|X_{\text{past}})}{P(X_{\text{int}})} \right] - \lambda \sum_{X_{\text{future}}} P(X_{\text{past}}|X_{\text{future}})P(X_{\text{future}}) \ln \left[\frac{P(X_{\text{int}}|X_{\text{future}})}{P(X_{\text{int}})} \right] - \mu(X_{\text{past}}). \tag{808}$$

To proceed, it would be useful to divide through by a factor of $P(X_{\text{past}})$, at which point we have

$$\ln \left[\frac{P(X_{\text{int}}|X_{\text{past}})}{P(X_{\text{int}})} \right] = \lambda \sum_{X_{\text{future}}} P(X_{\text{future}}|X_{\text{past}}) \ln \left[\frac{P(X_{\text{int}}|X_{\text{future}})}{P(X_{\text{int}})} \right] + \tilde{\mu}(X_{\text{past}}), \tag{809}$$

where $\tilde{\mu}(X_{\text{past}}) = \mu(X_{\text{past}})/P(X_{\text{past}})$. Further, since on the right we have a conditional distribution of X_{future} given X_{past} , it would be nice to rearrange the ratio inside the logarithm,

$$\frac{P(X_{\text{int}}|X_{\text{future}})}{P(X_{\text{int}})} = \frac{P(X_{\text{future}}|X_{\text{int}})}{P(X_{\text{future}})} = \frac{P(X_{\text{future}}|X_{\text{int}})}{P(X_{\text{future}}|X_{\text{past}})} \cdot \frac{P(X_{\text{future}}|X_{\text{past}})}{P(X_{\text{future}})}, \tag{810}$$

so that, when we substitute we find

$$\begin{aligned}
\ln \left[\frac{P(X_{\text{int}}|X_{\text{past}})}{P(X_{\text{int}})} \right] = & \lambda \sum_{X_{\text{future}}} P(X_{\text{future}}|X_{\text{past}}) \ln \left[\frac{P(X_{\text{future}}|X_{\text{int}})}{P(X_{\text{future}}|X_{\text{past}})} \right] \\
& + \lambda \sum_{X_{\text{future}}} P(X_{\text{future}}|X_{\text{past}}) \ln \left[\frac{P(X_{\text{future}}|X_{\text{past}})}{P(X_{\text{future}})} \right] + \tilde{\mu}(X_{\text{past}}). \tag{811}
\end{aligned}$$

We recognize the first term on the right as being the (negative) Kullback–Leibler divergence between the distribution of futures given the past, and the distribution

of futures given our representation X_{int} . Further, the second term depends only on X_{past} , and so can be absorbed into $\tilde{\mu}(X_{\text{past}})$. Thus, when the dust settles, we have

$$P(X_{\text{int}}|X_{\text{past}}) = \frac{P(X_{\text{int}})}{Z(X_{\text{past}}; \lambda)} \exp \left(- \lambda D_{KL} [P(X_{\text{future}}|X_{\text{past}}) || X_{\text{future}}|X_{\text{int}}] \right), \tag{812}$$

where $Z(X_{\text{past}}; \lambda)$ is a normalization constant. This isn't

a solution to our problem, but rather a self-consistent

equation that the solution has to satisfy. The problem we are solving is an example of selective compression, and the particular formulation of trading bits vs. bits has come to be called the “information bottleneck” problem.

Problem 155: Fill in all the details leading from Eq (807) to Eq (812).

We should think of Equation (812) as being like the result in Eq (794), but instead of adjusting an internal state in relation to the “potential” formed by the growth rate, here the effective potential is the (negative) Kullback–Leibler divergence, which measures the similarity between the distributions of futures given the actual past and given our compressed representation of the past. This means that if two past histories lead to similar distributions of futures, they should be mapped into the same value of X_{int} . This makes sense, since we are trying to throw away any information that doesn’t have predictive power. When λ is very large, differences in the expected future need to be very small before we are willing to ignore them, while at small λ it is more important that our description be compact, so we are willing to make coarser categories. As in rate–distortion theory, there is no single right answer, but rather a curve which defines the maximum amount of predictive information we can capture given that we are willing to write down a certain number of bits about the past, and along this curve there is a one parameter family of strategies for mapping our observations on the past into some internal representation X_{int} .

Problem 156: Predictive information and optimal filtering. Imagine that we observe a Gaussian stochastic process $[x(t)]$ that consists of a correlated signal $[s(t)]$ in a background of white noise $[\eta(t)]$, that is $x(t) = s(t) + \eta(t)$, where

$$\langle s(t)s(t') \rangle = \sigma^2 \exp(-|t - t'|/\tau_c) \quad (813)$$

$$\langle \eta(t)\eta(t') \rangle = \mathcal{N}_0 \delta(t - t'). \quad (814)$$

Recall (or see Section A.2) that the full probability distribution for the function $x(t)$ is

$$P[x(t)] = \frac{1}{Z} \exp \left[-\frac{1}{2} \int dt \int dt' x(t) K(t - t') x(t') \right], \quad (815)$$

where Z is a normalization constant.

(a.) Construct the kernel $K(\tau)$ explicitly. Be careful about the behavior near $\tau = 0$.

(b.) Break the data $x(t)$ into a past $X_{\text{past}} \equiv x(t < 0)$ and a future $X_{\text{future}} \equiv x(t > 0)$, relative to the time $t = 0$. Show that

$P[x(t)]$ can be rewritten so that the only term that mixes past and future is of the form

$$\left[\int_{-\infty}^0 dt g(-t)x(t) \right] \times \left[\int_0^{\infty} dt' g(t')x(t') \right], \quad (816)$$

where $g(t) = \exp(-t/\tau_0)$, with $\tau_0 = \tau_c(1 + \sigma^2\tau_c/\mathcal{N}_0)^{-1/2}$. More formally, if we define

$$z = \int_{-\infty}^0 dt g(-t)x(t), \quad (817)$$

show that

$$P(X_{\text{future}}|X_{\text{past}}) = P(X_{\text{future}}|z). \quad (818)$$

Explain why the optimal internal representation of the predictive information, X_{int} , can only depend on z .

(c.) Suppose that we are given the past data $x(t \leq 0)$, and instead of being asked to predict the future, you are asked to make the best estimate of the underlying signal $s(t = 0)$. [Connect back to problem in Chapter 1] Show that this optimal estimate is proportional to z .

As you just showed in the last problem, the optimal representation of predictive information is equivalent, at least in simple cases, to the separation of signals from




FIG. 146 A precursor of the information bottleneck problem, from Lee et al (1993). In one year of the Associated Press news reports, there are 64 nouns (X_{noun}) which appear as the direct object of the verb “to fire,” and these nouns are paired with 2147 distinct verbs (X_{verb}). Following the ideas in the text, imagine compressing the description of the nouns, $X_{\text{noun}} \rightarrow X_{\text{int}}$, while trying to preserve the information that the compressed description conveys about the verb which appears with the noun. That is, maximize $I(X_{\text{int}}; X_{\text{verb}})$ while holding $I(X_{\text{int}}; X_{\text{noun}})$ fixed. Here we show the solution to the problem when $I(X_{\text{int}}; X_{\text{noun}}) \approx 1$ bit supports two distinct values of X_{int} ; what we list are the nouns that map to the two values of X_{int} with high probability. We see that this classifies the nouns by their meaning, separating weapons (firing a missile) from job titles (firing a manager). Importantly, this is based only on the co-occurrence of the nouns with verbs in sentences; there is no supervisory signal which distinguishes the different senses of the verb.

noise. In Section IV.D we will see that extracting the predictive information from other kinds of time series is equivalent to learning the underlying parameters or rules that the data obey. In a somewhat more fanciful example, we can think of X_{past} as a word in a sentence, and X_{future} as the next word; then the mapping $X_{\text{past}} \rightarrow X_{\text{int}}$ is equivalent to making clusters of words. When λ is small, there are very few clusters, and they correspond very closely to parts of speech. As λ becomes larger, we start to discern categories of words that seem to have meaning. Indeed the first exercise of this sort was to choose not two successive words as past and future, but rather the noun and verb in the same sentence, and then the impression (still subjective) that the resulting clusters of nouns have similar meanings is even stronger, as seen in Fig 146. It is tempting to suggest that the optimal representation of predictive information is extracting “meaning” from the statistics of sentences. [perhaps explain that some people are horrified by this suggestion?]

[Maybe say something about Tagkopoulos et al (2008)? What about different ideas of predictive coding from Laughlin, Rao, ... ? Is this the place to show evidence (depending on how much we have!) that neurons provide efficient representations of predictive information, or does this go in the next section? What about a reminder that the rules of synaptic plasticity seem to know about causality, and hence might serve to build representations that favor predictive information?]

Let me try to pull the different arguments of this section together, even if imperfectly. What we really care about is how organisms can maximize some measure of performance—ultimately, their reproductive success—given access to some limited set of resources. Within any broad class of possible biological mechanisms, there is an optimum that divides the fitness/resources plane into possible and impossible regions, as in the upper right quadrant of Fig 147; evolutionary pressure drives organisms toward this boundary. But we have seen that, for any measure of fitness or adaptive value, achieving some criterion level of performance always requires some minimum number of bits; this is the content of rate–distortion theory. Thus there is a plane (in the upper left quadrant of Fig 147) of fitness vs. information, and again there is a curve that divides the possible from the impossible. Importantly, the information that an organism can use to gain a fitness advantage—even in the simple example of adjusting gene expression levels to match the availability of nutrients—is always predictive information, because the consequences of actions come after they are decided upon.

We know that bits are not free. In simple examples, such as the Gaussian channel in Section ??, the information that can be transmitted depends on the signal to noise ratio, and this in turn depends on the resources the organism can devote, whether we are counting action potentials or molecules. If we think about the bits that will

be used to direct an action, then there are many costs—the cost of acquiring the information, of representing the information, and the more obvious physical costs of carrying out the resulting actions, but we always can assign these costs to the symbols at the entrance to the communication channel. The channel capacity separates the information/resources plane into accessible and inaccessible regions, as in the lower right quadrant of Fig 147. Ideas about metabolically efficient neural codes [perhaps should be more explicit here?], for example, can be seen as efforts to calculate this curve in specific models. Of course the information we are talking about now is information that we actually collect, and this is information about the past. To close the connections among the different quantities, we need the information bottleneck, which tells us that—given the structure of the world we live in—having a certain number of bits of information about the future requires capturing some minimum number of bits about the past.

To summarize, if an organism wants to achieve a certain mean fitness, it needs a minimum number of bits of predictive power, and this requires collecting a minimum number of bits about the past, which in turn necessitates some minimum cost or available resources. Usually we




FIG. 147 Connecting the different optimization principles (Bialek et al 2007). Lines indicate curves of optimal performance, separating allowed from forbidden (hashed) regions of each quadrant. In the upper right quadrant is the biologically relevant notion of optimization, maximizing fitness or adaptive value at fixed resources. But actions that achieve a given level of adaptive value require a minimum number of bits, and since actions occur after plans these are bits about the future (upper left). On the other hand, the organism has to “pay” for bits, and hence there is a minimum resource costs for any representation of information (lower right). Finally, given some bits (necessarily obtained from observations on the past), there is some maximum number of bits of predictive power (lower left). To find a point on the biological optimum one can try to follow a path through the other three quadrants, as indicated by the arrows.

think of evolution as operating in the tradeoff between resources and fitness, but this has echoes in the other quadrants of Fig 147, where information theoretic bounds are at work. These connections provide a path whereby evolution can select for mechanisms that approach these bounds, even though evolution itself doesn't know about bits.

The connection between information and gambling goes back to Kelly (1956). Connections of these ideas to fitness in fluctuating environments are discussed by Bergstrom & Lachmann (2005), Kussell & Leibler (2005), and more generally by Rivoire & Leibler (2011). The specific case of persistence in bacteria has been explored by Balaban et al (2004) and Kussell et al (2005); for a review see Gefen & Balaban (2009). The analogy to rate–distortion theory, demonstrating a minimum number of bits required to achieve a criterion mean growth rate, is from Taylor et al (2007); for a treatment of rate–distortion theory itself, again see Cover & Thomas (1991), in the refs to Section IV.A. Although they didn't explicitly use the language of rate–distortion theory, Park and Levitt (1995) explored the compression of protein structures into a small set of local, discrete states, asking how the complexity of this representation related to its accuracy. For a first try at connecting information flow and embryonic pattern formation, see Tkačik et al (2008). The beautiful convection patterns in Fig 144 are from Bodenschatz et al (1991).

Balaban et al 2004: Bacterial persistence as a phenotypic switch. NQ Balaban, J Merrin, R Chait, L Kowalik & S Leibler, *Science* **305**, 1622–1625 (2004).

Bergstrom & Lachmann 2005: The fitness value of information. CT Bergstrom & M Lachmann, arXiv:q-bio.PE/0510007 (2005).

Bodenschatz et al 1991: Transitions between patterns in thermal convection. E Bodenschatz, JR de Bruyn, G Ahlers & DS Cannell, *Physical Review Letters* **67**, 3078–3081 (1991).

Gefen & Balaban 2009: The importance of being persistent: Heterogeneity of bacterial populations under antibiotic stress. O Gefen & NQ Balaban, *FEMS Microbiol Rev* **33**, 704–717 (2009).

Kelly 1956: A new interpretation of information rate. JL Kelly, Jr, *Bell Sys Tech J* **35**, 917–926 (1956).

Kussell et al 2005: Bacterial persistence: A model of survival in changing environments. EL Kussell, R Kishony, NQ Balaban & S Leibler, *Genetics* **169**, 1807–1814 (2005).

Kussell & Leibler 2005: Phenotypic diversity, population growth, and information in fluctuating environments. E Kussell & S Leibler, *Science* **309**, 2075–2078 (2005).

Park & Levitt 1995: The complexity and accuracy of discrete state models of protein structure. BH Park & M Levitt, *J Mol Biol* **249**, 493–507 (1995).

Rivoire & Leibler 2011: The value of information for populations in varying environments. O Rivoire & S Leibler, *J Stat Phys* **142**, 1124–1166 (2011); arXiv.org:1010.5092 [q-bio.PE] (2010).

Taylor et al 2007: Information and fitness. SF Taylor, N Tishby & W Bialek, arXiv:0712.4382 [q-bio.PE] (2007).

Tkačik et al 2008: Information flow and optimization in transcriptional regulation. G Tkačik, CG Callan Jr & W Bialek, *Proc Nat'l Acad Sci (USA)* **105**, 12265–12270 (2008).

The 'information bottleneck' was introduced by Tishby et al (1999). It has connections with statistical mechanics approaches to clustering (Rose et al 1990), and more immediate antecedents in the idea of clustering distributions of words (Pereira et al 1993). For more about predictive information, see Bialek et al (2001) and (2007). The possibility that even familiar programs of coordinated changes in bacterial gene expression may reflect (implicit) predictions is discussed by Tagkopoulos et al (2008). Energy efficiency in neural coding is discussed by Laughlin et al (1998) and by Balasubramanian et al (2001).

Balasubramanian et al 2001: Metabolically efficient information processing. V Balasubramanian, D Kimber & MJ Berry II, *Neural Comp* **13**, 799–815(2001).

Bialek et al 2001: Predictability, complexity and learning. W Bialek, I Nemenman & N Tishby, *Neural Comp* **13**, 2409–2463 (2001); arXiv:physics/0007070 (2000).

Bialek et al 2007: Efficient representation as a design principles for neural coding and computation. W Bialek, RR de Ruyter van Steveninck & N Tishby, arXiv:0712.4381 [q-bio.NC] (2007); see also *Proceedings of the International Symposium on Information Theory 2006*.

Laughlin et al 1998: The metabolic cost of neural information. SB Laughlin, RR de Ruyter van Steveninck & JC Anderson, *Nature Neurosci* **1**, 36–41 (1998).

Pereira et al 1993: Distributional clustering of English words. FC Pereira, N Tisby & L Lee, in *30th Annual Meeting of the Association for Computational Linguistics*, pp 183–190 (1993).

Rose et al 1990: Statistical mechanics and phase transitions in clustering. K Rose, E Gurewitz & GC Fox, *Phys Rev Lett* **65**, 945–948 (1990).

Tagkopoulos et al 2008: Predictive behavior within microbial genetic network. I Tagkopoulos, Y Liu & S Tavazoie, *Science* **320**, 1313–1317 (2008).

Tishby et al 1999: The information bottleneck method. N Tishby, FC Pereira & W Bialek, in *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing*, B Hajek & RS Sreenivas, eds, pp 368–377 (University of Illinois, 1999); physics/0004057 (2000).

C. Optimizing information flow

We have seen that organisms should care about bits—for every criterion level of performance that a system wants to achieve, there is a minimum number of bits that it needs. If bits are cheap, or easy to acquire, then this need for a minimum number of bits is true but not much of a constraint. On the other hand, if the physical constraints under which organisms operate imply severe limits on information transmission, then the minimum number of bits may approach the maximum number available, and strategies that maximize efficiency in this sense may be critical to biological function.

One of the central ideas in thinking about the efficiency with which bits can be collected and transmitted is that what we mean by efficient (and, in the extreme, optimal) depends on context, as indicated schematically in

Fig 148. In the top panel we see a typical sigmoidal input/output relation, which might describe the expression level of a gene vs. the concentration of a transcription factor, the probability of spiking in a neuron as a function of the intensity of the sensory stimulus, In the bottom panel we see different possibilities for the distributions out of which the input signals might be drawn. For the two distributions in blue, the input signals are confined to the saturated regions of the input/output relation, leaving the output almost always in the fully ‘off’ or ‘on’ states. In these situations, the output is always the same, and is unaffected by the changes in the input that actually occur with reasonable probability, and the system is essentially useless. More subtly, for the distribution in green, input signals are in the middle of the slope of the input/output relation, where the slope of the input/output relation is maximal, but the dynamic range of these variations is small, so that the variations in output are only a small fraction of what is possible, and these variations might well be obscured by any reasonable level of noise. Finally, for the distribution in red, the dynamic range of the likely inputs is just big enough to push the system through the full dynamic range of the input/output relation, generating large (maximal?) variations in the output.




FIG. 148 At top, an example of an input/output relation. At bottom, different possible probability distributions for the inputs. As described in the text, the blue and green distributions are poorly matched to the input/output relation, while the red distribution seems to be a better match.

While it is easy for everyone to agree that, in Fig 148, the blue and green distributions of inputs are poorly matched to the input/output relation, and the red dis-

tribution is well matched, it takes a little more courage (and courts more controversy) to make a precise mathematical statement about what constitutes a good match, or the “best” match. What we will try out as a definition of “best” is that outputs should provide as much information as possible about the inputs.

Let’s start with input x , chosen from a distribution $P_X(x)$, and assume that this is converted into one output y by a system that has an input/output relation $g(x)$ but also some added noise,

$$y = g(x) + \xi. \quad (819)$$

Notice that when we plot an input/output relation, as in Fig 148, we (implicitly) are referring to the *average* behavior of the system, since realistically there must be some level of noise and hence the input and output are related only probabilistically; we now make this explicit by adding the noise ξ . To keep things simple, let’s assume that this noise is Gaussian, with some variance σ^2 and as usual zero mean. In principle, the variance of the output noise could depend upon the value of the input, and this will be important below, so we’ll write $\sigma_y^2(x)$ to remind us that we are talking about the variance of the output (hence the subscript), but this may depend upon the input.

In order to compute the amount of information that y provides about x , we need various probability distributions. Specifically, we want to evaluate

$$\begin{aligned} I(y; x) &= \int dx \int dy P(x, y) \log_2 \left[\frac{P(x, y)}{P_X(x)P_Y(y)} \right] \\ &= \int dx \int dy P(x, y) \log_2 \left[\frac{P(y|x)}{P_Y(y)} \right]. \end{aligned} \quad (820)$$

It is the conditional distribution $P(y|x)$ that describes, in the most general setting, the probabilistic relationship between input and output. The overall distribution of outputs is given by

$$P_Y(y) = \int dx P(y|x)P_X(x). \quad (821)$$

With the hypothesis that the noise ξ is Gaussian, Eq (819) tells us that

$$P(y|x) = \frac{1}{\sqrt{2\pi\sigma_y^2(x)}} \exp \left[-\frac{(y - g(x))^2}{2\sigma_y^2(x)} \right]. \quad (822)$$

The information can be written (as usual) as the difference between two entropies,

$$\begin{aligned}
I(y; x) &= \int dx \int dy P(x, y) \log_2 \left[\frac{P(y|x)}{P_Y(y)} \right] \\
&= - \int dy P_Y(y) \log_2 P_Y(y) - \int dx P_X(x) \left[- \int dy P(y|x) \log_2 P(y|x) \right].
\end{aligned} \tag{824}$$

But the conditional distribution $P(y|x)$ is Gaussian, with variance $\sigma_y^2(x)$, so we can substitute for the conditional entropy from Eq (701) to give

$$I(y; x) = - \int dy P_Y(y) \log_2 P_Y(y) - \frac{1}{2 \ln 2} \int dx P_X(x) \ln[2\pi e \sigma_y^2(x)]. \tag{825}$$

The distribution of outputs $P_Y(y)$ is broadened by two effects. First, as x varies, the mean value of y changes. Second, even with x fixed, noise causes variations in y . But if the noise is small, the first effect should dominate, and this will simplify our problem. Formally,

$$P_Y(y) = \int dx P_X(x) P(y|x) = \int dx P_X(x) \frac{1}{\sqrt{2\pi\sigma_y^2(x)}} \exp \left[- \frac{(y - g(x))^2}{2\sigma_y^2(x)} \right] \tag{826}$$

$$= \int dz \left| \frac{dz}{dx} \right|^{-1} P_X(x = g^{-1}(z)) \frac{1}{\sqrt{2\pi\sigma_y^2(z)}} \exp \left[- \frac{(y - z)^2}{2\sigma_y^2(z)} \right], \tag{827}$$

where we have changed variables to $z = g(x)$, which is allowed if the input/output relation is monotonic. But now we can view the integral as an average over a distribution of z , and we know that if the noise is small we can always write

$$\int dz F(z) \frac{1}{\sqrt{2\pi\sigma_y^2(z)}} \exp \left[- \frac{(y - z)^2}{2\sigma_y^2(z)} \right] \approx F(z = y) + \frac{1}{2} \sigma_y^2(z = y) \frac{d^2 F(z)}{dz^2} \Big|_{z=y} + \dots, \tag{828}$$

for any function $F(z)$. Keeping just the leading term, at small noise levels we have

$$P_Y(y) \approx \left[\left| \frac{dz}{dx} \right|^{-1} P_X(x = g^{-1}(z)) \right]_{z=y}. \tag{829}$$

This looks complicated, but it's not. In fact it is the same as ignoring the noise all together and saying that there is some deterministic transformation from x to y , $y = g(x)$, in which case we must have

$$P_X(x) dx = P_Y(y) dy. \tag{830}$$

By the same reasoning, we can also view the variance $\sigma_y^2(x)$ as being a function not of the input x but rather of the output y , so we'll write $\sigma_y^2(y)$.

In the small noise approximation, then, the mutual information between x and y thus can be written as

$$\begin{aligned}
I(y; x) &\approx - \int dy P_Y(y) \log_2 P_Y(y) \\
&\quad - \frac{1}{2 \ln 2} \int dy P_Y(y) \ln[2\pi e \sigma_y^2(y)].
\end{aligned} \tag{831}$$

Now it's clear that, given the noise level, we can maximize the mutual information by varying the distribution of outputs $P_Y(y)$. Notice that we started with the problem of varying the distribution of inputs, but now things are formulated in terms of the distribution of outputs; Eq (830) tells us that these are equivalent in the low noise limit. To do the optimization correctly, however, we have to add a Lagrange multiplier that fixes the normalization of the distribution. Thus we are interested in the functional

$$\tilde{I} \equiv I(y; x) - \mu \int dy P_Y(y). \tag{832}$$

As usual, to optimize we set the derivative equal to zero:

Problem 157: Details of the small noise approximation, part one. Show that Eq (830) really is the same as Eq (829).

$$\left. \frac{\delta \tilde{I}}{\delta P_Y(y)} \right|_{P_Y(y)=P_{\text{opt}}(y)} = 0 \quad (833)$$

$$\Rightarrow 0 = -\frac{1}{\ln 2} [\ln P_{\text{opt}}(y) + 1] - \frac{1}{2 \ln 2} \ln[2\pi e \sigma_y^2(y)] - \mu \quad (834)$$

$$\ln P_{\text{opt}}(y) = -\frac{1}{2} \ln[2\pi e \sigma_y^2(y)] - (1 + \mu \ln 2) \quad (835)$$

$$P_{\text{opt}}(y) = \frac{1}{\sqrt{2\pi e \sigma_y^2(y)}} e^{-(1+\mu \ln 2)}. \quad (836)$$

We can write this more simply by gathering together the various constants,

$$P_{\text{opt}}(y) = \frac{1}{Z} \frac{1}{\sigma_y}, \quad (837)$$

where Z must be chosen so that the distribution is normalized, so

$$Z = \int \frac{dy}{\sigma_y}. \quad (838)$$

With this result for the optimal distribution, the mutual information is

$$I_{\text{opt}} = \log_2 \left[\frac{Z}{\sqrt{2\pi e}} \right]. \quad (839)$$

Problem 158: Extrema of the mutual information. Once again we need to check that we have found an optimum, rather than some other type of extremum, in the dependence of the mutual information on the distribution of outputs, Eq (831). You can do this explicitly by computing second (functional) derivatives, or by appealing to general convexity properties of the entropy. Notice that our ability to write the information so simply as a functional of the output distribution alone is a feature of the low noise approximation. More generally, we should view the mutual information as a functional of the input distribution $P(x)$ and the conditional distribution(s) $P(y|x)$. Show that, in this more general setting, once $P(y|x)$ is known, the mutual information has a well defined maximum as a functional of $P(x)$.

Problem 159: Details of the small noise approximation, part two. Carry out the small noise approximation to the next leading order in the noise level σ_y^2 . Step by step, you should find $P(y)$ and then an expression for the information $I(y;x)$. What can you say about the problem of optimizing $I(y;x)$ in this case?

The result for the optimal distribution of outputs, Eq (837), is telling us something sensible: we should use the different outputs y in inverse proportion to how noisy they are. Suppose, however, that the noise level is constant. Then what we find is that the distribution of outputs should be uniform. How can the system do this?

Recall that in the low noise limit, the relationship between input and output is nearly deterministic, so we have Eq (830), $P_Y(y)dy = P_X(x)dx$. But we also have that $y = g(x)$, in this approximation. If $P_Y(y)$ is uniform, this means that

$$P_Y(y) = \frac{1}{y_{\text{max}} - y_{\text{min}}}, \quad (840)$$

and hence

$$\frac{dy}{dx} = \frac{dg(x)}{dx} = (y_{\text{max}} - y_{\text{min}})P(x) \quad (841)$$

$$g(x) = (y_{\text{max}} - y_{\text{min}}) \int_{x_{\text{min}}}^x dx' P(x'). \quad (842)$$

Thus, in this simple limit, the optimal input/output relation is proportional to the cumulative probability distribution of the input signals.

Problem 160: How general is Eq (842)? We have derived Eq (842) by assuming that the noise is additive, Gaussian, small, and finally has a variance that is constant across the range of inputs or outputs. Show that you can relax the assumption of Gaussianity (while keeping the noise small, additive, and independent of inputs) and still obtain the same result for the optimal input/output relation.

Equation (842) makes clear that any theory which involves optimizing information transmission or efficiency of representation inevitably predicts that the input/output relation must be matched to the statistics of the inputs. Here the matching is simple: in the right units we could just read off the distribution of inputs by looking at the (differentiated) input/output relation. Although this is obviously an over-simplified problem, it is tempting to test the predictions, and this is exactly what Laughlin did in the context of the fly's visual system.

Laughlin built an electronic photodetector with aperture and spectral sensitivity matched to those of the fly retina, and used this to scan natural scenes, sampling the distribution of input light intensities $P(\mathcal{I})$ as it would




FIG. 149 Input/output relations of large monopolar cells compared with the prediction of Eq (842), from Laughlin (1981). Brief changes in light intensity relative to a mean background produce transient voltage changes in the LMCs (inset), and the peaks of these responses are taken as the cell's output. Normalized responses are compared to the cumulative probability distribution of light intensities, as described in the text.

appear at the input to these neurons. In parallel he characterized the second order neurons of the fly visual system—the large monopolar cells which receive direct synaptic input from the photoreceptors, and which we have seen before in [pointers!]—by measuring the peak voltage response to flashes of light. The agreement with Eq (842) was remarkable, as shown in Fig 149, especially when we remember that there are no free parameters. While there are obvious open questions, this is a really beautiful result that inspires us to take these ideas more seriously.

This simple model automatically carries some predictions about adaptation to overall light levels. If we live in a world with diffuse light sources that are not directly visible, then the intensity which reaches us at a point is the product of the effective brightness of the source and some local reflectances. As is it gets dark outside the reflectances don't change—these are material properties—and so we expect that the distribution $P(\mathcal{I})$ will look the same except for scaling. Equivalently, if we view the input as the log of the intensity, then to a good approximation $P(\log \mathcal{I})$ just shifts linearly along the $\log \mathcal{I}$ axis as mean light intensity goes up and down. But then the optimal input/output relation $g(\mathcal{I})$ would exhibit a similar invariant shape with shifts along the input axis when expressed as a function of $\log \mathcal{I}$, and this is in rough agreement with experiments on light/dark adaptation in a wide variety of visual neurons [show a figure that illustrates this!].

As I have emphasized before, the problems of signals, noise and information flow in the nervous system have

analogous within the biochemical and genetic machinery of single cells. For the simple problem of one input and one output, we can move beyond analogy and actually use the same equations to describe these very different biological systems.

Suppose that we have a single transcription factor that controls the expression of one target gene. Now we can think of the input x as the concentration of the transcription factor, and the output y as the expression level of the gene. As in Laughlin's discussion of the fly retina, we are (perhaps dangerously) ignoring dynamics. In the context of gene regulation this probably is best seen as a quasi-steady state approximation, in which the changes in transcription factor concentration are either slow or infrequent, so that the resulting gene expression level has a chance to find its appropriate steady level in response.

We have discussed the problems of noise in the control of gene expression in Section II.B, and a crucial feature of that discussion is that the noise levels cannot be constant. In the simplest case, we are counting molecules, and counting zero molecules allows for no variance, while counting the maximum number of molecules leaves lots of room for variation. For the problem at hand, this means—because of Eq (837)—that the distribution of outputs that maximizes information transmission can't be uniform. To find the form of the optimal distribution we need to recall some of our earlier discussion about noise.

We have identified (at least) three noise sources in the regulation of gene expression. One term is the shot noise in the synthesis and degradation of the mRNA or protein (output noise). The second is the randomness in the arrival of transcription factor molecules at their target site (input noise), and the third is from the kinetics of the




FIG. 150 Optimal distributions of (output) gene expression levels. As described in the text, we maximize the transmission of information from a single transcription factor to a single target gene. Different curves correspond to relative contributions from input and output noise, as in Eq (845).

‘switching’ events that occur on binding of the transcription factors. We have argued that cells can reduce the impact of this last term by proper choice of parameters, leaving two fundamental sources of noise. The shot noise generates a variance at the output proportional to the mean, while the random arrivals are equivalent to a fluctuation in input concentration $(\delta c/c)^2 \propto 1/c$. Putting these together we have [from the discussion leading to Eq (376)] the variance in the expression level

$$\sigma_g^2(c) = \alpha \bar{g}(c) + \frac{B}{c} \cdot \left| \frac{d\bar{g}(c)}{d \ln c} \right|^2, \quad (843)$$

where α and B are constants, and $\bar{g}(c)$ is the mean expression level as a function of the input transcription factor concentration c ; as usual we will normalize the measurements of expression levels so that the maximum $\bar{g}(c) = 1$. Finally, if we can assume that the input/output relation is well approximated by a Hill function,

$$\bar{g}(c) = \frac{c^n}{c^n + K^n}, \quad (844)$$

then we can write the variance as a function of the mean, as in Eq (376),

$$\sigma_g^2(\bar{g}) = \alpha \bar{g} + \beta \bar{g}^{2-1/n} (1 - \bar{g})^{2+1/n}. \quad (845)$$

The parameter $A = \beta/\alpha$ measure the relative importance of input and output noise; large A means that the input noise is dominant near the midpoint of the input/output relation.

In Figure 150 we see the results for the optimal distributions of expression levels, derived using the general result of Eq (837) with the noise variance from Eq (845). We hold the cooperativity fixed ($n = 5$) and consider what happens as we change the relative importance of the input and output noise (A). As long as output noise is dominant, $A < 1$, the optimal distribution is monotonically decreasing. If we take the results seriously, the distribution has a singularity as we approach zero expression level. There is no physical reason why this can’t happen, but we also can’t trust our calculation here since at some point the noise $\sigma \propto \sqrt{\bar{g}}$ will become larger than the mean as $\bar{g} \rightarrow 0$. Nonetheless, it’s clear that when output noise is dominant, the optimal distribution of expression levels is relatively featureless, biased toward low expression levels. The strength of this bias is considerable so that the probability of having more than half-maximal activation, $\int_{1/2}^1 dg P(g)$, is a bit less than 30%.

At larger values of A , where input noise is more important, the optimal distribution of expression levels becomes bimodal. This is especially interesting, because extreme bimodality corresponds to a simple on/off switch. Intuitively, true switch-like behavior runs counter to the idea that information transmission is being maximized:




FIG. 151 Distributions of Hunchback expression levels in the early fruit fly embryo (Tkačik et al 2008). In red, the distribution predicted by optimizing information transmission given the measured input/output relation and noise in the control of Hb by Bcd. In black, with error bars, the distribution measured experimentally.

we might expect that maximizing information transmission involves making extensive use of intermediate expression levels, while building a reliable switch means exactly the opposite, avoiding intermediate levels. In fact, few of classic examples of “genetic switches” are perfect, and here we see that maximizing information transmission can lead to relatively low probabilities of occupying intermediate levels, just depending on the structure of the noise in the system.

We can bring this theoretical discussion down to earth by considering a real system. As discussed in Section II.B, there are measurements on the input/output relation and noise level for the control of the *hunchback* gene by the transcription factor Bicoid in the early *Drosophila* embryo. If we take the formalism above seriously, we can use these measurements to predict, with no free parameters, the distribution of *hunchback* expression levels, which can also be extracted from the experiments. To do this correctly, we should go beyond the small noise approximation and solve the full optimization problem numerically; the results are shown in Fig 151.

Figure 151 is the direct analog of Laughlin’s result in the fly retina. As in that case, the agreement of theory and experiment is very good, and again it should be emphasized that there are no free parameters—these are not models we are fitting to data, but quantitative predictions from theory. One can go further, and show from the data that the actual amount of information⁸⁶ being transmitted from Bicoid to Hunchabck is 0.88 ± 0.09 of

⁸⁶ This is a good place to remember the technical difficulties involved in estimating information from finite samples of data. See Appendix A.9.




FIG. 152 Changes in information transmission from Bcd to Hb if we scale all noise variances by a factor F , from Tkačik et al (2008). This is equivalent to scaling all the numbers of molecules by a factor $1/F$. At each noise level we compute the maximum information transmission, as described in the text. Limiting behaviors in the small and large noise approximations are shown for reference. The real system (in red, with error bar) is in an intermediate regime, although close to the small noise limit.

the limit set by the measured noise levels. Thus, going back to the remarks in Section 1.5, we can see directly that the system is operating near its optimum. This optimum corresponds to significantly more than one bit, which means that intermediate expression levels, beyond an on/off switch, are being used reliably. Finally, since we understand how the absolute numbers of molecules influence the noise level in the system (see, again, Section II.B), we can compute that more bits would be very expensive—doubling the information would require twenty times as many molecules, as shown in Fig 152.

Problem 161: Information flow through calcium binding proteins. Many biological processes are regulated by calcium. Typically the regulatory process begins with calcium binding to a protein. In almost all cases, there are multiple binding sites, and these sites interact cooperatively. We'd like to understand something about the signals, noise and information flow in such regulatory systems; not much has been done in this area so this is a deliberately open ended problem. Consider the simple model shown in Fig 153. This is a dimeric protein with four states, corresponding to empty and filled Ca^{++} binding sites on each of the two monomers. The sites interact, since the rate of unbinding from one site depends on the occupancy of the other site.

(a.) Calculate the equilibrium probability of occupying each of the states in Fig 153. Use these results to plot the fraction of occupied binding sites as a function of the calcium concentration c . You should be able to choose units which eliminate all parameters except for the dimensionless constant F . Show that for $F = 1$ your results are equivalent to having two independent binding sites, and that the fraction of occupied sites becomes more strongly sigmoidal or switch-like as F becomes larger. Cooperativity means, in this

context, that the free energy change upon binding of a calcium ion to one site is increased by occupancy of the other site. Relate the parameter F to this free energy difference or interaction energy. Can interaction energies of just a few times $k_B T$ make a difference in the shape of the plot of occupancy vs. concentration? See also the discussion of cooperativity in Appendix A.4.

(b.) Suppose that we have N copies of this protein in the cell, all experiencing the same calcium concentration. Let the number of molecules with no bound calcium be n_0 , the number with one bound calcium be n_1 , and the number with two bound calcium be n_2 ; of course $\sum_j n_j = N$. Use your results from Problem 1 to calculate the mean values of each n_j and the covariance matrix $C_{jk} = \langle \delta n_j \delta n_k \rangle$. Verify that the determinant of the covariance matrix is zero in this formulation. Why is this true? Notice that we're only asking here about the fluctuations that you would see in a single snapshot of the molecules, not about the dynamics or spectrum of this noise.

(c.) It is widely assumed that in systems such as this, only the state with full occupancy of the binding sites is really “active.” In practice what this means is that the calcium binding protein is associated with some other protein, such as an enzyme, and the enzyme becomes active only when both Ca^{++} are bound. Thus, the output of the system is something proportional to n_2 . Calculate the change in the mean $\langle n_2 \rangle$ that results from a small change in calcium concentration $c \rightarrow c + \delta c$. Compare this with the variance $\langle (\delta n_2)^2 \rangle$ to compute a signal-to-noise ratio, or the equivalent noise level δc_{rms} in the calcium concentration itself. Plot your results. Again, you should be able to put everything into unitless form, leaving only the parameter F . Does making the system more switch-like by increasing F makes it more sensitive to small changes in concentration, as you might expect? Are there competing effects which could result in better performance at smaller F ?

(d.) Suppose that molecules with one bound calcium also are active. Then the output activity of the system is proportional to some mixture of n_1 and n_2 , which we can write as $A = (1-a)n_2 + an_1$; note that $a = 0$ brings us back to the case where only doubly-bound states are active. Compute sensitivity of the mean activity, $\partial \langle A \rangle / \partial c$ and the variance $\langle (\delta A)^2 \rangle$. If the system is operating at a particular calcium concentration c , can you lower the effective noise level

$$\delta c_{\text{rms}} \equiv \sqrt{\langle (\delta A)^2 \rangle} \left| \frac{\partial \langle A \rangle}{\partial c} \right|^{-1}$$




FIG. 153 Model of calcium binding to a dimeric protein. The rate at which calcium binds to each site, k_+ is assumed to be the same and independent of the occupancy of the other site. The unbinding rates, however, are different depending on whether the other site is empty (k_-) or filled (k_-/F).

by choosing $a \neq 0$? Can you lower the noise level at all calcium concentrations using the same value of a , or are there tradeoffs?

(e.) One way to think about the effective noise level δc_{rms} is that it sets a scale for the smallest concentration differences that can be detected. If we imagine that c can range from zero up to some maximum c_{max} , then it seems natural to say that number of different levels of concentration that can be distinguished is given by

$$N_{\text{levels}} = \int_0^{c_{\text{max}}} \frac{dc}{\delta c_{\text{rms}}(c)}, \quad (846)$$

where we note explicitly that the noise level depends on the background concentration. The number of distinguishable levels should translate into an information transmission (in bits) of $I \sim \log_2(N_{\text{levels}})$, and this is almost right in the limit that the noise is small. Show how a rigorous version of this argument can be constructed by analogy with the derivation of Eq's (838) and (839). Calculate I for the system discussed above. Does thinking about the information transmitted, rather than just the noise level, help you to decide whether there is a uniquely best mixture of activity from the singly- and doubly-bound states? What is the impact of the cooperativity (here captured by the parameter F) on the information transmission?

(f.) Some things to explore: Your results above suggest that, at least under some conditions, it would be useful if the system “reads out” some combination of the singly- and doubly-occupied states. Can you find hints in the literature of the predicted partial activation? For concreteness, focus on the case of calmodulin. Our discussion above is for snapshots of the molecules, so ‘noise’ just means the total variance. Suppose that the readout scheme effectively averages over a time longer than the times required for transitions among the different states. Then you need to compute the spectral density of the noise, and follow the path we discussed in the context of bacterial chemotaxis. Is there anything qualitatively new here, or just a change in details?

One of many questions left open in Laughlin’s original discussion is the time scale on which the matching should occur. One could imagine that there is a well defined distribution of input signals, stable on very long time scales, in which case the matching could occur through evolution. Another possibility is that the distribution is learned during the lifetime of the individual organism, perhaps largely during the development of the brain to adulthood. Finally one could think about mechanisms of adaptation that would allow neurons to adjust their input/output relations in real time, tracking changes in the input distribution. It seems likely that the correct answer is all of the above. But the last possibility, real time tracking of the input distribution, is interesting because it opens the possibility for new experimental tests.

We know that some level of real time matching occurs, as in the example of light and dark adaptation in the visual system. We can think of this as neurons adjusting their input/output relations to match the mean of the input distribution. The real question, then, is whether there is adaptation to the distribution, or just to the mean. Actually, there is also a question about the world we live in, which is whether there are other features of the distribution that change slowly enough to be worth tracking in this sense.




FIG. 154 This is a placeholder .. should replace with real data, e.g. sounds from the songbird colony. Intermittency in natural sounds. Top trace shows the alternating “loud” and “soft” period characteristic of natural sounds. Probability distribution of the instantaneous signal amplitude is far from Gaussian, having long, nearly exponential tails. Need to illustrate more clearly that these tails are removed by local variance normalization!

As an example, we know that many signals that reach our sensory systems come from distributions that have long tails (cf Fig 154). In some cases (e.g., in olfaction, where the signal—odorant concentration—is a passive tracer of a turbulent flow) there are clear physical reasons for these tails, and indeed it’s been an important theoretical physics problem to understand this behavior quantitatively. In most cases, the tails arise through some form of intermittency. Thus, we can think of the distribution of signals as being approximately Gaussian, but the variance of this Gaussian itself fluctuates; samples from the tail of the distribution arise in places where the variance is large. This scenario also holds for images of the natural world, so that there are regions of high variance and regions of low variance. The possibility of such “variance normalization” in images suggests that the visual system could code more efficiently by adapting to the local variance, in addition to the local mean (light and dark adaptation).

Adaptation to local variance, or more generally adaptation to input statistics beyond the mean, definitely happens at many stages of neural processing (Fig 155). The earliest experiments looked the responses of retinal ganglion cells to sudden changes in the variance of their inputs, and showed that there is a pattern very similar to what one sees with sudden changes in mean. More ambitious experiments on the motion-sensitive neurons in the fly visual system mapped the input/output relation when inputs were drawn from different distributions, and found that the input/output relation scales in proportion to the dynamic range of inputs, which is what




FIG. 155 At left, adaptation of retinal ganglion cells to sudden changes in the variance of light intensity (Smirnakis et al 1997). At right top, input/output relations for the fly motion-sensitive neuron H1 measured when inputs are drawn from different distributions (Brenner et al 2000). To be precise one has to define the input as a filtered version of the velocity, and the methods for determining these filters are discussed in Appendix A.7. At right bottom, the input/output relations collapse when expressed as a function of the stimulus in units of its standard deviation.

one expects from the matching principle if noise levels are small; it was also checked that the precise proportionality constant in the scaling relation served to maximize information transmission. Further, if you suddenly switch from one distribution to another, you can ‘catch’ the system using the wrong code and transmitting less information, but the adaptation to the new distribution is very fast, close to the limit set by the need to collect enough samples that you are sure there was a change. Related observations have been in many systems, from low level sensory neurons up to mammalian cortex. [Do we want to say more here? Maybe work on how such adaptation is a property of individual neurons, so it is a building block of neural computation? At least pointers to the fact that these effects are so fast that calling them “adaptation” raises some questions.]

I think the adaptation experiments are important because they give a whole new way of testing the ideas about matching between the input/output relation and the distribution of inputs—by changing the input distribution, if you believe the theory, we should drive changes in the input/output relation, and it seems that this works. Can we imagine a similar experiment in the genetic or biochemical systems? In truth, there are few cases (aside from embryonic development) where we have quantitative measurements on the distributions of inputs under moderately natural conditions. If we change the distribution, then for the case of gene regulation one imagines that input/output relations could change in re-

sponse only on evolutionary time scales, but at least for bacteria such evolutionary experiments are now quite feasible. Certainly there are models for network evolution that use information theoretic quantities as a surrogate for fitness, and these models are generating interesting predictions, as shown in Fig [include a figure from Francois & Siggia simulations]. It would be exciting to see laboratory evolution experiments that are the analog of the neural experiments in Fig 155.

So far the discussion is about one input and one output, in single genes or neurons. Almost all the really interesting systems, however, involve populations or networks of these elements. Indeed, one of the earliest ideas about optimizing information transmission in neural coding is that interactions among neighboring neurons in the retina serve to reduce the redundancy of the signals that they transmit, thus making better use of their capacity.

To get a feeling for how redundancy reduction works, consider a system in which there are N receptor cells that produce signals x_i , and these feed into a layer of N output neurons that take linear combinations of their inputs and add noise, so that the outputs of the system are

$$y_i = \sum_j W_{ij} x_j + \eta_i, \quad (847)$$

and shown schematically in Fig 847. In the simplest case the noise will be Gaussian and independent in each output neuron, $\langle \eta_i \eta_j \rangle = \delta_{ij} \sigma^2$. Let’s also assume, again for simplicity, that the distribution of the x s is also Gaussian, with zero mean and a covariance matrix $\langle x_i x_j \rangle = C_{ij}$. Then following the arguments in Section IV.A, the information that the outputs provide about the inputs is

$$I(\vec{y}; \vec{x}) = \frac{1}{2} \text{Tr} \log_2 \left(\mathbf{1} + \frac{1}{\sigma^2} W C W^T \right), \quad (848)$$

where W^T denotes the transpose of the matrix W . We can chose the matrix W , which defines the “receptive fields” of the output neurons [point back to first discussion of receptive fields; check!] to maximize the information, but we need a constraint, since otherwise the answer is always to make W larger so we can overwhelm the noise. A natural constraint, then, is to fix the overall dynamic range of the output signal,

$$\sum_i \langle y_i^2 \rangle = \text{Tr} (W C W^T) + N \sigma^2. \quad (849)$$

But if we go into a basis where $W C W^T$ is diagonal, then the information becomes

$$I(\vec{y}; \vec{x}) = \frac{1}{2} \sum_{\mu} \log_2 \left(1 + \frac{\Lambda_{\mu}}{\sigma^2} \right), \quad (850)$$

where the Λ_{μ} are the eigenvalues of the matrix $W C W^T$, and the constraint is that the sum of these eigenvalues

must be constant. Then it is clear from the convexity of the logarithm that the best we can do is to have all of the eigenvalues be equal, which means that WCW^T is proportional to the unit matrix. But (leaving aside the contribution from the noise), WCW^T is the correlation matrix of the output signals. Thus, in this simple model, we maximize information transmission by removing all of the correlations in the input, and making the outputs independent of one another.

Problem 162: Convexity and equalization. Show explicitly that if we want to maximize

$$I = \frac{1}{2} \sum_{\mu} \log_2(1 + \tilde{\Lambda}_{\mu}), \quad (851)$$

subject to the constraint

$$\sum_{\mu} \tilde{\Lambda}_{\mu} = C, \quad (852)$$

then the solution is to have all the $\{\tilde{\Lambda}_{\mu}\}$ be equal, that is $\tilde{\Lambda}_{\mu} = \Lambda_0$. [I'd like to get the students to think more about the implications of this ...]

In the retina, we expect that correlations, and perhaps also the transformations from input to output, are translation invariant. Thus if the receptor cell i is at position \mathbf{r}_i , perhaps on a lattice, and the output neurons are on the same lattice, we should have

$$C_{ij} = C(\mathbf{r}_i - \mathbf{r}_j), \quad (853)$$

$$W_{ij} = W(\mathbf{r}_i - \mathbf{r}_j). \quad (854)$$




FIG. 156 A schematic network, after Eq (847). The x_j provide inputs to the y_i , with weights W_{ij} . All connections are present, but the connections from x_3 are highlighted.




FIG. 157 Cross-sections through the optimal matrices W_{ij} in the problem with noise, from Atick & Redlich (1990). The correlation function is assumed to be exponential, $C_{ij} \propto \exp(-|i - j|/\xi)$, with $\xi = 50$, much longer than the range of interactions shown here. At high SNR, the solution looks like a differentiator, which decorrelates the signals, while at low SNR the solution integrates to suppress noise.

Then the condition for independence at the outputs becomes

$$\delta_{ij} \propto \sum_{km} W_{ik} C_{km} W_{jm} \quad (855)$$

$$= \sum_{km} W(\mathbf{r}_i - \mathbf{r}_k) C(\mathbf{r}_k - \mathbf{r}_m) W(\mathbf{r}_j - \mathbf{r}_m). \quad (856)$$

We approximate the sums as integrals, so that

$$\delta_{ij} \approx \int d^2r' \int d^2r'' W(\mathbf{r}_i - \mathbf{r}') C(\mathbf{r}' - \mathbf{r}'') W(\mathbf{r}_j - \mathbf{r}'') \quad (857)$$

$$= \int \frac{d^2k}{(2\pi)^2} |\tilde{W}(\mathbf{k})|^2 S(\mathbf{k}) e^{i\mathbf{k} \cdot (\mathbf{r}_i - \mathbf{r}_j)}, \quad (858)$$

where $\tilde{W}(\mathbf{k})$ is the Fourier transform of $W(\mathbf{r})$, and we identify the Fourier transform of the correlation function $C(\mathbf{r})$ as the power spectrum $S(\mathbf{k})$. To satisfy this condition $|\tilde{W}(\mathbf{k})|^2 S(\mathbf{k})$ must be constant, independent of \mathbf{k} , and if $W(\mathbf{r})$ is symmetric in space this means that

$$\tilde{W}(\mathbf{k}) \propto \frac{1}{\sqrt{S(\mathbf{k})}}. \quad (859)$$

We expect that the power spectrum of correlations in the inputs to the retina fall off at high frequencies, which means that the optimal weights W have the form of a filter which does the opposite, attenuating the low frequencies and enhancing high frequencies. In fact, experiments show that the power spectrum of contrast in natural scenes is scale invariant, so that $S(\mathbf{k}) \propto |\mathbf{k}|^{-\alpha}$, with the exponent α close to 2. Then the optimal weights W

should actually vanish as $\mathbf{k} \rightarrow 0$, which means that the output of the retina should be insensitive to spatially uniform illumination; on the other hand, the output should overemphasize gradients or edges. By including the effects of noise (as in the problem below), one can see a crossover to spatial averaging at low SNR; see Fig 157. Qualitatively this is all correct: at high signal-to-noise ratios we can see this enhancement of edges not just in the responses of retinal ganglion cells but also in our perception, through the phenomenon of Mach bands [figures?], and this spatial differentiation gives way to integration as we lower the light levels and hence the SNR.

Problem 163: Redundancy reduction vs noise reduction. Equation (859) suggests that at large \mathbf{k} , where the power spectrum of input signals should be small, the weight in transferring these signals to the output should be large. This can't be completely right, since we expect that at very high (spatial) frequencies, signals will be lost in a background of noise. Go back to the start of this analysis and assume that the signals x_i already have a little bit of noise attached to them (as with photon shot noise in vision) so that

$$y_i = \sum_j W_{ij}(x_j + \xi_j) + \eta_i, \quad (860)$$

where everything is as before but $\langle \xi_i \xi_j \rangle = \delta_{ij} \sigma_0^2$. Follow the outline above and derive the form of the weights W_{ij} that optimize information transmission at fixed output variance. Verify that as $\sigma_0 \rightarrow 0$ you recover the simple picture in which the optimal W_{ij} serve to remove correlations. Show, in contrast, that as σ_0 becomes large, the optimal solution involves averaging over multiple inputs to beat down the noise.

Problem 164: Information available at the retina. Give a problem that takes the students through the calculation in Ruderman & Bialek (1994), showing that with reasonable assumptions natural scenes provide only ~ 1 bit per cone in the fovea.

[Do we want to talk about coding/whitening in the time domain, maybe the results on filtering at the receptor/LMC synapse? Could argue by analogy with spatial whitening, give a problem to work out details.]

Maybe a simpler example of these ideas is provided by color processing. Roughly speaking, at one point in space our retina takes three samples, corresponding to the signals in the three different cones. These three signals are correlated, both because the absorption spectra of the pigments in the different cones overlap and because the reflectance spectra of the objects around us are rather smooth functions of wavelength.⁸⁷ By analogy with what we have seen thus far, if the retina is under pressure to maximize information transmission then

it should send these signals to the brain in some decorrelated form. Early guesses about the form of the correlations among the different cone signals suggested that the three decorrelated signals would correspond roughly to the sum of all the inputs (the total light intensity, ignoring color), an approximately “red minus green” signal, and a “blue minus yellow” signal. In fact it is known that neurons throughout the visual system follow this pattern of “opponent” color processing [feels like there should be something about experiments demonstrating color opponency, but I don't know how quantitatively one can make comparisons, so ...?].

To do a more quantitative analysis one has to get away from traditional color photography, because (for example) the three channels in a CCD camera don't have wavelength sensitivities that correspond exactly to that of our cones. Instead one can take hyperspectral images, essentially measuring the spectrum of light at each point in the scene, and then construct the expected signals that will be seen by each cone, known the absorption spectra of the three cone pigments. This analysis shows, quite remarkably, that the rough intuition about opponent processing is nearly exact, with the decorrelated signals being almost perfect integer combinations of the cone signals: if the three cone signals are \mathcal{L} , \mathcal{M} and \mathcal{S} for the long, medium and short wavelengths, then the decorrelated signals are $\ell = \mathcal{L} + \mathcal{M} + \mathcal{S}$ (light intensity), $\alpha = \mathcal{L} + \mathcal{M} - 2\mathcal{S}$ (blue minus yellow), and $\beta = \mathcal{L} - \mathcal{M}$ (red minus green), where the coefficients are unity with an accuracy of $\sim 1\%$. Further, this linear transformation serves to generate truly independent signals, even though the underlying distributions are not Gaussian; see Fig 158. These very clean results come from a delicate interplay between the statistical structure of the world and the properties of our visual pigments. I don't know how accurately the coefficients in opponent color processing have been measured, but this is a striking prediction that certainly captures the qualitative behavior of the system and deserves to be tested more quantitatively.

In the example of Fig 157, the optimal weights for transforming receptor signals into neural output correspond to a “center-surround” structure in which an output neuron at point \mathbf{r} gives a positive weight to the receptor cell at point \mathbf{r} , and a negative weight to its neighbors. Alternatively, we can think that all weights from receptors to neurons are positive, and the output neurons inhibit one another before sending their signals on to the brain. In our retinae things are complicated, because the transformation from photoreceptors to ganglion cells involves several intermediate cells, but in some simpler creatures such as the horseshoe crab the picture of “lateral inhibition” seems to be correct, and indeed the horseshoe crab was the first retina in which receptive fields were measured. Lateral inhibition is thought to be a general neural mechanism for “sharpening” the responses to stimuli that vary across an array of neu-

⁸⁷ In fact this is the same effect. The reflectance properties of most naturally occurring objects in our terrestrial environment are determined by the absorption spectra of organic pigments, and these tend to be broad; see Section [**] and Appendix [**].




FIG. 158 Statistical structure of color images, from Ruderman et al (1998). Top left shows a color image of one scene analyzed by hyperspectral imaging. From the raw data, one constructs the signals \mathcal{L} , \mathcal{M} and \mathcal{S} corresponding to the (log) photon capture rates by each of the types of cones, and then rotates into the basis defined by ℓ (upper right), α (lower left), and β (lower right), as described in the text; images are shown in these three projections after thresholding for clarity. The three images are uncorrelated.

rons, and we have seen that this sharpening is essential in decorrelating signals and enhancing the efficiency of information transmission. Could there be an analog of this for the transmission of information through genetic or biochemical networks? If we go back to the case of Bicoid regulating the expression of Hunchback, we know that this is just one piece of a larger network in which the primary morphogen Bicoid feeds into a collection of gap genes, which in turn interact with one another. Because transcription factors tend to be either activators or repressors, in the absence of any other effects all of the target genes would have correlated expression levels and hence provide redundant data about the concentration of the input. This redundancy can be removed by lateral inhibition, and that is what we see in the gap gene network (Fig 159). The challenge is to take this quantitative analogy and turn it into a quantitative theory.

The representations of data constructed by the nervous system might be efficient in the sense we have considered here, but they have a more obvious feature—they are built from discrete action potentials or spikes. If we look with some reasonably fine time resolution, $\Delta\tau < 10$ ms, then since the average spike rates are less than 100 spikes/s, at any moment the typical neuron is silent. In this sense, the code is “sparse.” It is this sparseness which, among other things, makes it possible to decode spike trains using linear filters, as in Eq (729). Spikes are expensive, requiring substantial energy expenditure, and perhaps it is this cost which drives the brain toward the construction of sparse representations.




FIG. 159 The gap gene network in the *Drosophila* embryo. **Need to check and see how much of this has been discussed already, although this is a nice place to put this ...**

If we take the idea of linear reconstruction seriously, then if the sensory input is $s(t)$ —for example sound pressure as a function of time in the auditory system—we would like to have a family of neurons labeled by μ that spike at times $\{t_i^\mu\}$ such that

$$s_{\text{est}}(t) = \sum_{\mu} \sum_i f_{\mu}(t - t_i^\mu) \quad (861)$$

is as close as possible to the true signal. Notice that in this system the input is $s(t)$ and the output is the set of spikes $\{t_i^\mu\}$. If we imagine adjusting the input/output relations, the mapping $s(t) \rightarrow \{t_i^\mu\}$ will change, perhaps in complicated ways. But suppose we knew the functions $f_{\mu}(\tau)$. Then there would be “best times” t_i^μ for each spike so that the match between $s_{\text{est}}(t)$ and $s(t)$ is as close as possible. We could imagine searching through some large space of input/output relations to find one that puts the spikes at these best times, or we could use the times themselves as our description of the input/output relation. Conversely, if we knew the spike times, we could adjust the filters $f_{\mu}(\tau)$, as in our previous discussions. Can we do both problems, subject to a constraint on the total number of spikes? This is hard, but by slightly softening the problem—allowing each term in Eq (861) to have a varying amplitude—it becomes tractable. **[Can we say something about whether the softening makes much difference in the end? Need to ask Lewicki for details.]**

Figure 160 shows the results of this approach applied to a small population of neurons “trained” to provide an efficient representation of natural sounds. There are several interesting features in these results. First, the filters $f_{\mu}(\tau)$ are localized in time; although they are “tuned” to particular frequencies, they are more like a wavelet than a Fourier representation, with support over a window of time that scales inversely with the characteristic fre-

quency. The filters also have a very asymmetric shape, with a sharp attack and a slower decay. If we look through measurements of the impulse responses of neurons emerging from the mammalian ear, we see exactly these structures, and one can even find cells that overlay the predicted filters almost perfectly.⁸⁸ Importantly, these structures are lost if one tries to build representations of very different sound ensembles.

By allowing for different total numbers of spikes, or limiting the time resolution with which the spikes are placed, one can construct codes of different qualities. For these different codes it is relatively easy to put an upper bound on the entropy of the spike trains, and to measure the errors between $s_{\text{est}}(t)$ and the true $s(t)$; putting these together one obtains the rate–distortion curve for




FIG. 160 Ingredients for an efficient representation of natural sounds, as in Eq (861), from Smith & Lewicki (2006). At left, the functions $f_{\mu}(\tau)$ in red, compared with the impulse responses of single neurons in the cat auditory nerve in blue; grey scale bars are 5 ms long. All these filters are band pass, so that their Fourier transforms $\hat{f}_{\mu}(\omega)$ have maximum magnitude at some characteristic frequency and fall to half maximal over some bandwidth. At right, a scatter plot of bandwidths vs characteristic frequencies for the filters (red) and auditory neurons (small blue dots); filters trained on different ensembles (black circles and green triangles) have very different behavior.

⁸⁸ One needs to be careful here. The impulse responses of the neurons are measured by the reverse correlation method (see Appendix A.7) which, ideally, extracts a filter characteristic of the *encoding* of sounds into spikes. In contrast, the filters $f_{\mu}(\tau)$ are characteristic of the *decoding* process. One can circumvent this problem by using reverse correlation to analyze the model code, with almost identical results. This suggests that the model, at least, is operating in a regime where the coding and decoding filters are similar. This happens exactly in the limit where all spikes are statistically independent from one another, so that there is no redundancy. Thus, the search for efficient codes may also drive the emergence of simplicity in decoding.

this family of codes. Applied to ensembles of human speech, the results are comparable to or better than conventional coding schemes. It does indeed seem that nature has found an efficient class of codes, not just in abstract terms.

Is there more we want to say here? What about predictive information? I'd love to say something about this, if we know enough

Laughlin's classic paper on matching input/output relations to the distribution of inputs still is very much worth reading, thirty years later (Laughlin 1981). The corresponding analysis for a genetic regulatory element is by Tkačik et al (2008a), with more theoretical exploration in Tkačik et al (2008b). The literature on information transmission in biochemical and genetic networks is growing rapidly; for examples see Ziv et al (2007), Mugler et al (2008), Yu et al (2008) and Tostvein & ten Wolde (2009). For a detailed model of calcium signaling via the protein calmodulin (of relevance to Problem **), see Pepke et al (2010).

Laughlin 1981: A simple coding procedure enhances a neuron's information capacity. SB Laughlin, *Z Naturforsch* **36c**, 910–912 (1981).

Mugler et al 2008: Form, function and information processing in small stochastic biological networks. A Mugler, E Ziv, I Nemenman & CH Wiggins, *IET Sys Biol* (2008).

Pepke et al 2010: A dynamic model of interactions of Ca^{2+} , calmodulin, and catalytic subunits of Ca^{2+} /calmodulin-dependent protein kinase II. S Pepke, T Kinzer–Ursem, S Mihalas & MB Kennedy, *PLoS Comp Biol* **6**, e1000675 (2010).

Tkačik et al 2008a: Information flow and optimization in transcriptional regulation. G Tkačik, CG Callan Jr & W Bialek, *Proc Nat'l Acad Sci (USA)* **105**, 12265–12270 (2008).

Tkačik et al 2008b: Information capacity of genetic regulatory elements. G Tkačik, CG Callan Jr & W Bialek, *Phys Rev E* **78**, 011910 (2008).

Tostvein & ten Wolde 2009: Mutual information between input and output trajectories of biochemical networks. F Tostvein & PR ten Wolde, *Phys Rev Lett* **102**, 21801 (2009).

Yu et al 2008: Negative feedback that improves information transmission in yeast signalling. RC Yu, CG Pesce, A Colman–Lerner, L Lok, D Pincus, E Serra, M Holl, K Benjamin, A Gordon & R Brent, *Nature* **456**, 755–761 (2008).

Ziv et al 2007: Optimal signal processing in small stochastic biochemical networks. E Ziv, I Nemenman & CH Wiggins, *PLoS* **2**, e1007 (2007).

The idea of matching input/output relations to the statistics of inputs had a big impact on neuroscience, in particular driving the exploration of input statistics under natural conditions. An early paper in this direction was by Field (1987), who noted that the power spectra of natural images were approximately scale invariant. Natural images are strongly nonGaussian, so we expect that scaling should mean much more than an appropriately shaped power spectrum, and this is true (Ruderman & Bialek 1994, Ruderman 1994); exploration of these statistical structures beyond the Gaussian approximation led to the ideas of variance normalization, and the prediction of adaptation to the local variance. Well before these analyses there was a substantial body of work on “contrast gain control” at various levels of the visual system; see, for example, Shapley & Victor (1981). The work on image statistics prompted a more explicit search for adaptation to the distribution

of visual inputs beyond the mean light intensity (Smirnakis et al 1997). Brenner et al (2000) describe experiments mapping the input/output relations of the fly's motion-sensitive visual neurons when inputs are drawn from different distributions, demonstrating that this adaptation served to optimize information transmission, and Fairhall et al (2001) explored the dynamics of this process, showing that one could "catch" the system using the wrong code and transmitting less information.

Brenner et al 2000: Adaptive rescaling optimizes information transmission. N Brenner, W Bialek & R de Ruyter van Steveninck, *Neuron* **26**, 695–702 (2000).

Fairhall et al 2001: Efficiency and ambiguity in an adaptive neural code. AL Fairhall, GD Lewen, W Bialek & RR de Ruyter van Steveninck, *Nature* **412**, 787–792 (2001).

Field 1987: Relations between the statistics of natural images and the response properties of cortical cells. DJ Field, *J Opt Soc Am A* **4**, 2379–2394 (1987).

Ruderman 1994: The statistics of natural images. DL Ruderman, *Network* **5**, 517–548 (1994).

Ruderman & Bialek 1994: Statistics of natural images: Scaling in the woods. DL Ruderman & W Bialek, *Phys Rev Lett* **73**, 814–817 (1994).

Shapley & Victor 1981: How the contrast gain control modifies the frequency responses of cat retinal ganglion cells. RM Shapley & JD Victor, *J Physiol (Lond)* **318**, 161–179 (1981).

Smirnakis et al 1997: Adaptation of retinal processing to image contrast and spatial scale. S Smirnakis, MJ Berry II, DK Warland, W Bialek & M Meister, *Nature* **386**, 69–73 (1997).

Adaptation to the distribution of inputs has now been reported in many neural systems: in the song bird (Nagel & Doupe 2006) and mammalian (Dean et al 2005, 2006) auditory systems, in the visual cortex (Sharpee et al 2006) and in the somatosensory system (Maravall et al 2007). For a review, including the connections to information transmission, see Wark et al (2007). The retina offers an accessible model system in which to explore the mechanisms of such statistical adaptation, and these mechanisms seem quite rich and diverse (Rieke 2001, Kim & Rieke 2001, Baccus & Meister 2002, Kim & Rieke 2003). **Most recent things from Fairhall & Rieke on the time constants of adaptation in retina. Look at recent work showing these effects in response to current injection in single neurons. Also say something about whether it is all too fast to be "real" adaptation, and whether this matters?** For ideas about the adaptation/evolution of biochemical and genetic networks, see Francois et al (2007) and Francois & Siggia (2010).

Baccus & Meister 2002: Fast and slow contrast adaptation in retinal circuitry. SA Baccus & M Meister, *Neuron* **36**, 900–919 (2002).

Dean et al 2005: Neural population coding of sound level adapts to stimulus statistics. I Dean, NS Harper & D McAlpine, *Nature Neurosci* **8**, 1684–1689 (2005).

Dean et al 2006: Rapid neural adaptation to sound level statistics. I Dean, BL Robinson, NS Harper & D McAlpine, *J Neurosci* **28**, 6430–6438 (2006).

Francois et al 2007: A case study of evolutionary computation of biochemical adaptation. P Francois, V Hakim & ED Siggia, *Mol Syst Biol* **3**, 154 (2007).

Francois & Siggia 2010: Predicting embryonic patterning using mutual entropy fitness and in silico evolution. P Francois & ED Siggia, *Development* **137**, 2385–2395. (2010)

Kim & Rieke 2001: Temporal contrast adaptation in the input and output signals of salamander retinal ganglion cells. KJ Kim & F Rieke, *J Neurosci* **21**, 287–299 (2001).

Kim & Rieke 2003: Slow Na⁺ inactivation and variance adaptation in salamander retinal ganglion cells. KJ Kim & F Rieke, *J Neurosci* **23**, 1506–1516 (2003).

Maravall et al 2007: Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. M Maravall, RS Petersen, AL Fairhall, E Arabzadeh & ME Diamond, *PLoS Biology* **5**, e19 (2007).

Nagel & Doupe 2006: Temporal processing and adaptation in the songbird auditory forebrain. KI Nagel & AJ Doupe, *Neuron* **51**, 845–859 (2006).

Rieke 2001: Temporal contrast adaptation in salamander bipolar cells. F Rieke, *J Neurosci* **21**, 9445–9454 (2001).

Sharpee et al 2006: Adaptive filtering enhances information transmission in visual cortex. TO Sharpee, H Sugihara, AV Kurgansky, SP Rebik, MP Stryker & KD Miller, *Nature* **439**, 936–942 (2006).

Wark et al 2007: Sensory adaptation. B Wark, BN Lundstrom & A Fairhall, *Curr Opin Neurobiol* **17**, 423–429 (2007).

The ideas about efficient coding in the population of retinal ganglion cells go back to Barlow (1959, 1961), in delightfully original papers that I still find very inspiring. The precise mathematical formulation of these ideas took until Atick & Redlich (1990); van Hateren (1992) worked out essentially the same principles with invertebrate rather than vertebrate retinas in mind, and there are important precursors in the work of Srinivasan et al (1982) and Snyder et al on compound eye design (cited in Section I.A). The possibility that opponent color processing is an example of efficient coding was suggested by Buchsbaum & Gottschalk (1983), and the full analysis based on hyperspectral images is due to Ruderman et al (1998). An attempt to derive gene regulatory networks that optimize information transmission is described Tkačik et al (2009) and Walczak et al (2010).

Atick & Redlich 1990: Toward a theory of early visual processing. JJ Atick & AN Redlich, *Neural Comp* **2**, 308–320 (1990).

Barlow 1959: Sensory mechanisms, the reduction of redundancy, and intelligence. HB Barlow, in *Proceedings of the Symposium on the Mechanization of Thought Processes, volume 2*, DV Blake & AM Utley, eds, pp 537–574 (HM Stationery Office, London, 1959).

Barlow 1961: Possible principles underlying the transformation of sensory messages. HB Barlow, in *Sensory Communication*, W Rosenblith, ed, pp 217–234 (MIT Press, Cambridge, 1961).

Buchsbaum & Gottschalk 1983: Trichromacy, opponent colour coding and optimum colour information transmission in the retina. G Buchsbaum & A Gottschalk, *Proc R Soc Lond B* **220**, 89–113 (1983).

van Hateren 1992: Real and optimal neural images in early vision. JH van Hateren, *Nature* **360**, 68–70 (1992).

Ruderman et al 1998: Statistics of cone responses to natural images: Implications for visual coding. DL Ruderman, TW Cronin & C-C Chiao, *J Opt Soc Am A* **15**, 2036–2045 (1998).

Srinivasan et al 1982: Predictive coding: A fresh view of inhibition in the retina. MV Srinivasan, SB Laughlin & A Dubs, *Proc R Soc Lond B* **216**, 427–459 (1982).

Tkačik et al 2009: Optimizing information flow in small genetic networks. I. G Tkačik, AM Walczak & W Bialek, *Phys Rev E* **80**, 031920 (2009).

Walczak et al 2010: Optimizing information flow in small genetic networks. II: Feed-forward interaction. AM Walczak, G Tkačik & W Bialek, *Phys Rev E* **81**, 041905 (2010).

Perhaps the most obvious evidence that the energy costs of spiking are significant is the fact that functional magnetic resonance imaging (fMRI) of the brain actually works: what one "sees" in these experiments are the changes in blood oxygenation that reflect the metabolic load associated with neural activity (Ogawa

et al 1990, 1992). The importance of fMRI has been one of the stimuli for a more detailed accounting of the energy budget of the brain (Atwell & Laughlin 2001, Raichle & Gusnard 2002). The description of sparse/efficient coding in the auditory system is based on the work of Lewicki (2002) and colleagues (Smith & Lewicki 2005, 2006). This grows out of earlier ideas of Lewicki & Sejnowski (2000). For an overview of sparse coding and spikes, see Olshausen (2002) [more detailed refs to these ideas in the context of visual cortex].

Atwell & Laughlin 2001: An energy budget for signaling in the grey matter of the brain. D Atwell & SB Laughlin, *J Cereb Blood Flow & Metab* **21**, 1133–1145 (2001).

Lewicki 2002: Efficient coding of natural sounds. MS Lewicki, *Nature Neurosci* **5**, 356–363 (2002).

Lewicki & Sejnowski 2000: Learning overcomplete representations. MS Lewicki & TJ Sejnowski, *Neural Comp* **12**, 337–365 (2000).

Ogawa et al 1990: Brain magnetic resonance imaging with contrast dependent on blood oxygenation. S Ogawa, TM Lee, AR Kay & DW Tank, *Proc Nat'l Acad Sci (USA)* **87**, 9868–9872 (1990).

Ogawa et al 1992: Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. S Ogawa, DW Tank, R Menon, JM Ellerman, SG Kim, H Merkle & K Ugurbil, *Proc Nat'l Acad Sci (USA)* **89**, 5951–5955 (1992).

Olshausen 2002: Sparse codes and spikes. BA Olshausen, *Probabilistic models of the brain: Perception and neural function*, RPN Rao BA Olshausen & MS Lewicki, eds. pp 257–272 (MIT Press, Cambridge, 2002).

Raichle & Gusnard 2002: Appraising the brain's energy budget. *Proc Nat'l Acad Sci (USA)* **99**, 10237–10239 (2002).

Smith & Lewicki 2005: Efficient coding of time-relative structure using spikes. EC Smith & MS Lewicki, *Neural Comp* **17**, 19–45 (2005).

Smith & Lewicki 2006: Efficient auditory coding. EC Smith & MS Lewicki, *Nature* **439**, 978–982 (2006).




FIG. 161 Fitting to polynomials. We have a collection of data points (black circles), $\{x_n, y_n\}$, and we try to fit these data with polynomials of different degree K , with $K = 1, 5, 20$. We see that as the degree of the polynomial—what we think of intuitively as the complexity of our model—increases, we can get closer to the data points, but at the same time we are introducing wild fluctuations which seem unlikely to be correct. In fact, $K = 5$ is the correct answer, since the data points were generated by choosing the x_n at random, evaluating some fixed fifth-order polynomial, and then adding noise. To claim that we understand how to learn, we have to find a principled way of convincing ourselves that it's better to keep the poorer fit with the simpler model.

in our physics lab classes (see Fig 161 for a reminder), and even this simple example introduces us to many deep issues. First, data usually come with some level of noise, and because of this any model really is (at least implicitly) a model of the probability distribution out of which the data are being drawn, rather than just a functional relationship. Indeed, one could argue that the general problem is always the problem of learning such distributions, and any rigid or deterministic rules emerge as a limit in which the noise becomes small or is beaten down by a large number of observations. The second point is that we would like to compare different models, often with different numbers of parameters. We have an intuition that simpler models are better, and we want to make this intuition precise—is it just a subjective preference, or is the search for simplicity something we can ground in more basic principles? A related point is that where the classical curve fitting exercises involve models with a limited number of parameters, we might want to go beyond this restriction and consider the possibility that the data are described by functions that are merely 'smooth' to some degree. Finally, we would like to quan-

D. Gathering information and making models

The world around us, thankfully, is a rather structured place. Whether we are doing a careful experiment in the laboratory or taking a walk through the woods, the signals that arrive at our brains are far from random noise; there seem to be some underlying regularities or rules. Surely one task that all organisms must face is the learning or extraction of these rules and regularities, making models of the world, either explicitly or implicitly. In this section, we will explore how learning and making models is related to the general problem of efficient representation.

Perhaps the simplest example of learning a rule is fitting a function to data—we believe in advance that the rule belongs to a class of possible rules that can be parameterized, and as we collect data we learn the values of the parameters. This is something we all learned about

tify how much we are learning—and how much *can* be learned—about the underlying rules given a limited set of data. If there are limits to how much we can learn, is it possible that biology has constructed learning machines which are efficient in some absolute sense, pushing up against these limits? So, let's plunge in ...

Imagine that we observe two streams of data x and y , or equivalently a stream of pairs $(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)$. Assume that we know in advance that the x 's are drawn independently and at random from a distribution $P(x)$, while the y 's are noisy versions of some function acting on x ,

$$y_n = f(x_n; \boldsymbol{\alpha}) + \eta_n, \quad (862)$$

where $f(x; \boldsymbol{\alpha})$ is one function from a class of functions parameterized by $\boldsymbol{\alpha} \equiv \{\alpha_1, \dots, \alpha_K\}$ and η_n is noise, which for simplicity we will assume is Gaussian with known variance σ^2 . We can even start with a *very* simple case, where the function class is just a linear combination of basis functions, so that

$$f(x; \boldsymbol{\alpha}) = \sum_{\mu=1}^K \alpha_{\mu} \phi_{\mu}(x). \quad (863)$$

The usual problem is to estimate, from N pairs $\{x_i, y_i\}$, the values of the parameters $\boldsymbol{\alpha}$; in favorable cases such as this we might even be able to find an effective regression formula. Probably you were taught that the way to do this is to compute χ^2 ,

$$\chi^2 = \sum_n \left| y_n - f(x_n; \boldsymbol{\alpha}) \right|^2, \quad (864)$$

and then minimize to find the correct parameters $\boldsymbol{\alpha}$. You may or may not have been taught *why* this is the right thing to do, and this is what we would like to understand here.

If we assume that our model, Eq (862), is correct, what is the probability that we observe the data points $\{x_n, y_n\}$? Let's start by asking about the locations of the points x_n where we get samples of the functional relationship between x and y . In the standard examples of curve fitting, the examples are given to us and there is nothing more to say; thus, we might as well assume that the points x_n are chosen randomly and independently out of some distribution $P(x)$, perhaps just the uniform distribution on some interval. One might ask if there is a good choice for the next x_{n+1} , perhaps a point that will give us the maximal information about the underlying parameters $\boldsymbol{\alpha}$. This is the problem faced in the design of experiments—how do we choose what to measure given what we already know?—but let's leave this aside for the moment.

If we assume that the points $\{x_n\}$ are chosen out of some distribution, then conveniently our model in Eq (862) is a statement about the conditional probability distribution of y_n given x_n . Specifically, y_n is a Gaussian random variable with a mean value of $f(x_n; \boldsymbol{\alpha})$ and a variance of σ^2 , so that

$$P(y_n|x_n, \boldsymbol{\alpha}) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{(y_n - f(x_n; \boldsymbol{\alpha}))^2}{2\sigma^2} \right]. \quad (865)$$

By hypothesis, the noise on every point is independent, which means that

$$P(\{y_n\}|\{x_n\}, \boldsymbol{\alpha}) = \prod_{n=1}^N P(y_n|x_n, \boldsymbol{\alpha}). \quad (866)$$

Now we can put things together to write the probability of the data given the parameters of the underlying model,

$$P(\{x_n, y_n\}|\boldsymbol{\alpha}) = \left[\prod_{n=1}^N P(y_n|x_n, \boldsymbol{\alpha}) \right] \times \left[\prod_n P(x_n) \right] \quad (867)$$

$$= \left[\prod_n P(x_n) \right] \prod_{n=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{(y_n - f(x_n; \boldsymbol{\alpha}))^2}{2\sigma^2} \right] \quad (868)$$

$$= \exp \left[\sum_{n=1}^N \ln P(x_n) - \frac{N}{2} \ln(2\pi\sigma^2) - \frac{\chi^2}{2\sigma^2} \right] \quad (869)$$

where we identify χ^2 from Eq (864). Notice that the only place where the parameters appear is in χ^2 , and $P \propto e^{-\chi^2/2\sigma^2}$. Thus, finding parameters which minimize χ^2 also serves to maximize the probability that our model

could have given rise to the data. This sounds like a good thing to do, and certainly maximizing the probability of the data (usually called “maximum likelihood”) feels more fundamental than minimizing χ^2 . But what are we

really accomplishing by maximizing P ?

We recall from Section IV.A that the entropy is the expectation value of $-\log P$, and that it is possible to encode signals so that the amount of “space” required to specify each signal uniquely is on average equal to the entropy. In such optimal encodings, each possible signal s drawn from $P(s)$ can be encoded in a space of $-\log_2 P(s)$ bits. Thus, any model probability distribution implicitly defines a scheme for coding signals that are drawn from that distribution, so if we make sure that our data have high probability in the distribution (small values of $-\log P$) then we also are making sure that our code or representation of these data is compact. What this means is that good old fashioned curve fitting really is all about finding efficient representations of data, which is the same principle that we discussed in the previous section in contexts ranging from the regulation of gene expression to neural coding. To be clear, in the earlier discussion we took for granted some physical or resource constraint (e.g., the noise level or limited number of molecules) and tried to transmit as much information as possible. Here we do the problem the other way, searching for a representation of the data that will

require the minimum set of resources.

If we follow this notion of efficient representation a little further we can do better than just maximizing χ^2 . The claim that a model provides a code for the data is not complete, because at some point we have to represent our knowledge of the model itself. One idea is to do this explicitly—estimate how accurately you know each of the parameters, and then count how many bits you’ll need to write down the parameters to that accuracy and add this to the length of your code. Another idea is more implicit—you don’t really know the parameters, all you do is estimate them from the data, so it’s not so obvious that you should separate coding the data from coding the parameters, although this might emerge as an approximation. In this view what we should do is to integrate over all possible values of the parameters, weighted by some prior knowledge, and thus compute the probability that our data could have arisen from the *class* of models we are considering.

To carry out this program of computing the total probability of the data given the model class we need to do the integral

$$P(\{x_i, y_i\}|\text{class}) = \int d^K \alpha P(\alpha) P[\{x_i, y_i\}|\alpha] \quad (870)$$

$$= \int d^K \alpha P(\alpha) \exp \left[-\frac{N}{2} \ln(2\pi\sigma^2) - \frac{1}{2\sigma^2} \chi^2(\alpha; \{x_i, y_i\}) \right] \left[\prod_n P(x_n) \right], \quad (871)$$

where $P(\alpha)$ is the a priori distribution of parameters, maybe just a uniform distribution on some bounded region. But remember that χ^2 as we have defined it is a sum over data points, which means it (typically) will be proportional to N . Thus, at large N we are doing an integral in which the exponential has terms proportional to N —and so we should use a saddle point approximation. To implement this approximation let’s write

$$P(\{x_i, y_i\}|\text{class}) = \exp \left[-\frac{N}{2} \ln(2\pi\sigma^2) \right] \left[\prod_n P(x_n) \right] \int d^K \alpha e^{-Nf(\alpha)}, \quad (872)$$

where the effective “energy per data point” is

$$f(\alpha) = \frac{1}{2N\sigma^2} \chi^2(\alpha; \{x_i, y_i\}) - \frac{1}{N} \ln P(\alpha) \quad (873)$$

The saddle point approximation is that

$$\int d^K \alpha e^{-Nf(\alpha)} \approx e^{-Nf(\alpha^*)} (2\pi)^{K/2} \exp \left[-\frac{1}{2} \ln \det(N\mathcal{H}) \right] \quad (874)$$

where α^* is the value of α at which $f(\alpha)$ minimized, and

the Hessian \mathcal{H} is the matrix of second derivatives of f at this point,

$$\mathcal{H}_{\mu\nu} = \left. \frac{\partial^2 f(\alpha)}{\partial \alpha_\mu \partial \alpha_\nu} \right|_{\alpha=\alpha^*}. \quad (875)$$

At large N , $f(\alpha)$ is dominated by χ^2 , so α^* must be close to the point where χ^2 is minimized. Putting the pieces together, we have

$$-\ln P(\{x_i, y_i\}|\text{class}) \approx \frac{N}{2} \ln(2\pi\sigma^2) - \sum_{i=1}^N \ln P(x_i) + \frac{\chi_{\min}^2}{2\sigma^2} + \ln P(\alpha^*) - \frac{K}{2} \ln 2\pi\sigma^2 + \frac{1}{2} \ln \det(N\mathcal{H}). \quad (876)$$

Note that \mathcal{H} is a $K \times K$ matrix, and so $\det(N\mathcal{H}) = N^K \det(\mathcal{H})$. This allows us to group together terms based on their N dependence,

$$-\ln P(\{x_i, y_i\}|\text{class}) \approx - \sum_{i=1}^N \ln P(x_i) + \frac{\chi_{\min}^2}{2\sigma^2} + \frac{N}{2} \ln(2\pi\sigma^2) + \frac{K}{2} \ln N + \dots, \quad (877)$$

where the first three terms are $\propto N$, and the terms \dots (including things we have neglected in the saddle point approximation) are constant or decreasing as $N \rightarrow \infty$. Again, the negative log probability measures the length of the shortest code for $\{x_i, y_i\}$ that can be generated given the class of models.

In Equation (877), the first term averages to N times the entropy of the distribution $P(x)$, which makes sense since by hypothesis the x 's are being chosen at random. The second and third terms are as before, the length of the code required to describe the deviations of the data from the predictions of the best fit model; this also grows in proportion to N . The fourth term must be related to coding our knowledge of the model itself, since it is proportional to the number of parameters. We can understand the $(1/2) \ln N$ because each parameter is determined to an accuracy of $\sim 1/\sqrt{N}$, as in Fig 162, so if we start with a parameter space of size ~ 1 there is a reduction in volume by a factor of \sqrt{N} and hence a decrease in entropy (gain in information) by $(1/2) \ln N$. Finally, the terms \dots don't grow with N .

Problem 165: Deriving the code length in a class of models. Fill in the details leading to Eq (877). Find an explicit form for the terms \dots , and show that they do not grow with N . What assumptions do you need to make about the prior distribution $P(\alpha)$ in order to make this work?

What is crucial about the term $(K/2) \ln N$ is that it depends explicitly on the number of parameters. In general we expect that by considering models with more parameters we can get a better fit to the data, which means that χ^2 can be reduced by considering more complex model classes. But we know intuitively that this has to stop—we don't want to use arbitrarily complex models, even if they do provide a good fit to what we have seen. It is attractive, then, that if we look for the shortest code which can be generated by a class of models, there is an implicit penalty or coding cost for increased complexity. It is interesting from a physicist's point of view that this term




FIG. 162 Confidence limits on the estimation of mean and variance for a Gaussian distribution. In several independent experiments, we choose $N = 10$ (blue), $N = 100$ (green), or $N = 1000$ points out of a Gaussian distribution with zero mean and unit variance. We estimate the mean and variance from the data in the usual way, and draw error ellipses on the parameters that should contain 95% of the weight. We see that the linear dimensions of these ellipses shrink by $\sim 1/\sqrt{10}$ as N increase by a factor of 10. The (log) area inside the ellipses measures the entropy of our uncertainty in parameters, and decreases in this area correspond to gains in information.

emerges essentially from consideration of phase space or volumes in model space. It thus is an entropy-like quantity in its own right, and the selection of the best model class could be thought of as a tradeoff between this entropy and the “energy” measured by χ^2 , a view to which we return below.

Thus Eq (877) tells that we have a *natural* penalty for the complexity of our model. While this term is linear in the number of parameters, it is only logarithmic in the number of data points. In contrast, χ_{\min}^2 decreases with the number of parameters and is linear in the num-

ber of data point. In this way, the penalty for complexity becomes (relatively) less important the more data we gather: if we have only a few data points then although we could lower χ^2 by fitting every wiggle, the phase space factor pushes us away from this solution toward simpler models; if, however, the wiggles are consistent as we collect more data, then this factor becomes less important and we can move to the more complex models.

To see that these words really correspond to a quantitative theory, we have to generate a data set and go through the process of fitting via minimization of the ‘code length’ in Eq (877). For simplicity let’s consider polynomial functions. We can pick a polynomial by choosing coefficients a_μ at random, say in the interval $-1 < a < 1$, where

$$f(x) = \sum_{\mu=0}^{K_{\text{true}}} a_\mu x^\mu. \quad (878)$$

We’ll confine our attention to the range $-5 < x < 5$; in this range the function $f(x)$ has some overall dynamic range (measured, for example, by its variance over this interval), and we’ll assume the noise variance σ^2 is one percent of this ‘signal’ variance. Then we can generate points according to

$$y_n = f(x_n) + \eta_n, \quad (879)$$

and try to fit. Fitting to any polynomial of degree K by minimizing χ^2 is a standard exercise, and in this way we find $\chi_{\text{min}}^2(K)$. Then we can find the value of K that minimizes the total code length in Eq (877); this last step is just a competition between $\chi_{\text{min}}^2(K)/\sigma^2$ and $(K + 1) \ln N$. The results of this exercise are shown in Fig 163.

What we see in Fig 163 is that our qualitative description of the competition between complexity and goodness of fit really works. First we note that with a large number of data points, minimizing the code length zeroes in on the correct order of the underlying polynomial ($K \rightarrow K_{\text{true}}$), despite the presence of noise that one could ‘fit’ using more complex models. Next, we see that for smaller numbers of data points, the shortest code is biased toward simpler models. In the limit that we only have a handful of data points, the shortest code is often a straight line ($K = 1$). Put another way, we start with a bias toward simple models, and only as we uncover more data can we support the adding of greater complexity.




FIG. 163 Fitting to polynomials, part two. Choose the coefficients of a polynomial with degree 10 at random, and then choose points at random in the interval $-5 < x < 5$; there is added noise [as in Eq (862)] with a standard deviation set to be 1/10 of the overall dynamic range of the function $f(x)$. We then try to fit polynomials of order K , and find the value of K that minimizes the ‘code length’ in Eq (877). Since the result depends both on the particular value of the polynomial coefficients and on the particular points x_n that we happen to sample, we choose 500 examples and look at the mean (points) and standard deviation (error bars) across this ensemble of examples. Although the optimal order of the polynomial in any given example is, of course, an integer, fractional values arise from averaging over many examples.

verifying that higher order polynomials always give “better” fits in the sense of smaller χ^2 .

(c.) Notice that χ_{min}^2 is a function of K and N , but also a function of the particular points $\{x_i, y_i\}$ you have “observed” in the experiment and of the particular parameters $\{a_\mu\}$ that specify the real function you are trying to learn. When you choose a different set of parameters and test points $\{x_i\}$, from the same distribution, how different is the minimum “energy per data point” $\epsilon_{\text{min}} = \chi_{\text{min}}^2/(N\sigma^2)$ as a function of K ? What happens to this variability as N gets larger?

(d.) Perhaps the most important thing is to verify that minimizing the code length really does control the complexity of the fit, selecting a nontrivial optimum K . Convince yourself that, as in Fig 163, the optimal K is small than K_{true} for small data sets, and approaches K_{true} as you analyze larger data sets.

Problem 166: Fitting and complexity. Generate a version of Fig 163 for yourself, doing a simulation which follows the steps outlined in the text. If you do this in MATLAB, you’ll find the command `polyfit` to be useful. Some things to keep in mind:

(a.) Start with a small version of the problem, e.g. fitting to $N = 20$ data points.

(b.) Plot some of your intermediate results, just to get a feeling for what is going on. In particular, plot χ_{min}^2 as a function of K ,

There are many reasons to prefer simpler models, and certainly the idea that we entertain more complex models only as we collect more data is in accord with our sense of how we understand the world. But all of this can seem a little soft and squishy. Indeed, given the evident complexity of life and the world around us, one might start

to suspect that the preference for simple models is not an objective principle, but rather a subjective choice made by humans—and more often by scientists than by humans in general.⁸⁹ Even some technical discussions leave this impression of subjectivity, suggesting that while there must be a tradeoff between goodness of fit and complexity, the structure of this tradeoff is something that we are free to choose, perhaps inventing a new “penalty for complexity” tuned to the details of each problem. As physicists we are raised to be suspicious of overly complex models, but again this preference for simplicity is often couched in (surprisingly) soft words about the elegance or brevity of the equations that describe the model. What we have seen here is that all of this can be made much more precise.

The power of information theory in this context is that, by consistently measuring code lengths in bits we don’t have to discuss our ‘preference for simplicity’ as a separate principle from goodness of fit. Deviations from the model (badness of fit) and the complexity of the model both add bits to the overall code length, and the relative contributions are calculable with no adjustable constants. The absence of unknown constants is important, since if we had to specify weights for the different terms we would once again inject subjectivity into the discussion of just how much we care about simplicity. Instead, we have one principle (search for the most compact description) and everything else follows. In particular, what follows is that limited experience (small N) biases us toward simpler models, while as we accumulate more experiences (ultimately, as $N \rightarrow \infty$) we can admit more complex descriptions of the world.

This is a very satisfying picture, and I am inclined to say that we can declare victory—we understand what we are doing when we make models, why simple models are preferable, and how the support for more complex models emerges. Nonetheless, there are several loose ends, and I’m not sure that I know how to tie them all up.

The first and most obvious problem is that our discussion makes sense as long as we specify in advance a class of models, and more seriously a hierarchy of such classes with increasing complexity. It’s not at all obvious how to do this. Worse yet, plausible but wrong ways of doing this can lead to weird results, for example if we have a function well described by a Fourier series with just a few terms, but we try fitting polynomials. Simplicity and complexity have meaning as code lengths only if we have a defined ensemble of possibilities to choose from, in much the same way that Shannon’s original discussion of

the information gained on hearing the answer to a question (Section IV.A) starts with the assumption that we know the distribution out of which answers will be drawn.

A second, and perhaps related, problem is that we are discussing models with a finite number of parameters. It might seem more natural, for example, to imagine that the relationship between x and y is just some smooth function, not necessarily describable with a finite number of parameters; that is, $f(x)$ should live in a function space and not in a finite dimensional vector space. Now we have to specify a prior distribution not on the parameters, as with $P(\alpha)$ above, but on the functions themselves $P[f(x)]$. The simplest version of this problem is not with functional relations but just with probability distributions: suppose that we observe a set of points x_1, x_2, \dots, x_N , which we assume are drawn randomly and independently out of a distribution $Q(x)$; how do we estimate Q ? If the distribution we are looking for belongs to a family with a finite number of parameters, we proceed as before, but if all we know is that $Q(x)$ is a smooth function then we have to specify a prior probability distribution on this space of distributions. From a physicist’s point of view, probability distributions on such function spaces are just scalar field theories, and one can carry a fair bit of technology over to do real computations. The lesson from these computations is that, with some reasonable priors to implement what we mean by “smooth,” everything works as it does in the case of finite parameters, but the prior does matter.

Problem 167: Taming the singularities. The basic problem in trying to learn a continuous probability distribution is to explain why, having observed a set of points x_1, x_2, \dots, x_N , we shouldn’t just guess that the distribution is of the form

$$Q(x) \sim \frac{1}{N} \sum_{i=1}^N \delta(x - x_i), \quad (880)$$

which of course generates precisely the data we have observed with maximal (infinite!) probability density. We all know that this is the wrong answer, and the role of priors on the space of distributions is to express this knowledge. A very different approach to taming the singularities is sometimes called Kernel density estimation, in which we search for a probability distribution in the form

$$Q(x) = \frac{1}{K} \sum_{j=1}^K \frac{1}{\ell} F\left(\frac{x - y_j}{\ell}\right), \quad (881)$$

where ℓ is again a characteristic length scale, $F(z)$ is some ‘blob-like’ function, and the y_j are the centers of the blobs; F is normalized so that $\int dz F(z) = 1$. For concreteness let

$$F(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}. \quad (882)$$

If we let $K = N$ (generally not such a good idea), then it should be clear that the model which generates the data with the highest probability is one in which the kernel centers are on top of the data points, $y_i = x_i$ for all i . It should also be clear that this probability of the data increases for smaller ℓ , diverging as $\ell \rightarrow 0$. But we

⁸⁹ One could add that even among scientists, physicists have a special affinity for simple models, often to the point of being the punchline in jokes, as in “... consider the case of the spherical horse.”

know that, to get control over complexity, we should compute the *total* probability of generating the data in this class of model. In this case the parameters of the model are the kernel centers $\{y_i\}$. Assume that everything happens in a box, so that $0 < x < L$, and similarly for $\{y_i\}$; by translation invariance the prior on the y s should be flat in this box. Calculate the total probability that this class of models generates the data in the limit $\ell \rightarrow 0$. Is the answer finite? If so, this means that the phase space factors are just strong enough to compensate for the ‘goodness of fit’ and prevent anything from diverging in this limit. Can you find any other approximations that allow you to say anything about the optimal value of ℓ ?

Quite generally, when we compute the total probability

$$P(\{x_i\}|\text{model class}) = \left[\prod_{i=1}^N Q_{\text{true}}(x_i) \right] \int DQ P[Q(x)] \prod_{i=1}^N \left[\frac{Q(x_i)}{Q_{\text{true}}(x_i)} \right]. \quad (884)$$

We can collect the product into an exponential,

$$P(\{x_i\}|\text{model class}) = \left[\prod_{i=1}^N Q_{\text{true}}(x_i) \right] \int DQ P[Q(x)] \exp \left[N \frac{1}{N} \sum_{i=1}^N \ln \left(\frac{Q(x_i)}{Q_{\text{true}}(x_i)} \right) \right], \quad (885)$$

and we recognize that the average over data points x_i approaches, at large N , an average over the true distribution,

$$P(\{x_i\}|\text{model class}) \rightarrow \left[\prod_{i=1}^N Q_{\text{true}}(x_i) \right] \int DQ P[Q(x)] \exp \left[N \int dx Q_{\text{true}}(x) \ln \left(\frac{Q(x)}{Q_{\text{true}}(x)} \right) \right] \quad (886)$$

$$= \left[\prod_{i=1}^N Q_{\text{true}}(x_i) \right] \int d\epsilon \mathcal{N}(\epsilon) e^{-N\epsilon}, \quad (887)$$

where

$$\epsilon = \int dx Q_{\text{true}}(x) \ln \left(\frac{Q_{\text{true}}(x)}{Q(x)} \right) \quad (888)$$

counts the (weighted) volume in model space that is at KL divergence ϵ away from the right answer. Now ϵ , which is a “goodness of fit” between the model and the data, can be thought of as an energy, while the (log) volume in model space is an entropy, $\mathcal{N}(\epsilon) = e^{S(\epsilon)}$. If we imagine the the model space has a finite but large dimensionality K , then we expect that the entropy will be extensive,

that a model can generate data, we are doing integrals like

$$P(\{x_i\}|\text{model class}) = \int DQ P[Q(x)] \prod_{i=1}^N Q(x_i), \quad (883)$$

where $P[Q(x)]$ is the probability distribution function(al) on the space of distributions. It embodies all our prior knowledge, in whatever form—that the distribution can be described by a few parameters, or merely that it is smooth in some sense. To understand what is happening in this integral, it is useful to measure possible distributions $Q(x)$ relative to the true distribution $Q_{\text{true}}(x)$,

is the Kullback–Leibler divergence between the distribution $Q(x)$ and the true distribution, and

$$\mathcal{N}(\epsilon) = \int DQ P[Q(x)] \delta \left[\epsilon - \int dx Q_{\text{true}}(x) \ln \left(\frac{Q_{\text{true}}(x)}{Q(x)} \right) \right] \quad (889)$$

$S(\epsilon) = K s(\epsilon)$. So, when the dust settles,

$$P(\{x_i\}|\text{model class}) \propto \int d\epsilon \exp \left[-N \left(\epsilon - \frac{K}{N} s(\epsilon) \right) \right]. \quad (890)$$

Thus, at large N , the integral is dominated by the minimum of the free energy density, $f = \epsilon - Ts(\epsilon)$, where the role of temperature is played by $T = K/N$. This calculation makes explicit the idea that learning really is statistical mechanics in the space of models, and that see-

ing more examples is like lowering the temperature, ‘cooling’ the system into an ordered state around the right answer. Depending on space of possible models, and hence the function $s(\epsilon)$, there can be phase transitions—a sudden jump, as we collect more examples, from wandering around in model space to having a compelling fit to the data.

What would it mean to have a phase transition in learning? As we accumulate more examples, we are lowering the effective temperature in the equivalent statistical mechanics problem. At first this doesn’t do very much, in the same way that lowering the temperature of water from 100 °C to 30 °C doesn’t do very much. But, at some point, a relatively small change in the number of examples we have seen produces a huge change in the distribution over models, freezing into a small volume surrounding the correct answer. This would be something like the subjective “aha!” experience, where we suddenly seem to understand something or master a skill after a very period of experience or training. Although we have all (I hope) experienced this phenomenon, it is not so easy to study quantitatively, and so I think we have no idea whether the statistical mechanics approach to learning provides a useful guide to understanding this effect.

It is interesting to look at the history of studies in animal learning in the light of these results. Already in the 1920s and 30s it was clear that, at certain tasks, animals could exhibit “sudden” rather than gradual learning. Although this was well before Hebb, and decades before the observation of changes in synaptic strength driven by the correlation between pre- and post-synaptic neurons (see Section [\[point back to previous chapter; be sure it’s there!\]](#)), there was a general view that learning relied upon statistical association, and thus should be a continuous process. Thus there was a question whether sudden learning represents a new mechanism, beyond associative processes. The mapping of learning onto a statistical mechanics problem reminds us that when there are many degrees of freedom, continuous dynamics can have nearly discontinuous consequences.

Before leaving the image of energy/entropy competition behind, we should note a caveat. In getting to Eq (890), we have first allowed N to become very large, so that averages over samples can be replaced by averages over the underlying distribution, and then used the resulting formulae with finite N to say something about how learning proceeds. Evidently this is dangerous. It also was controversial when it first emerged, since the results seemed to conflict with an approach by computer scientists which emphasized bounds on the learning curve. To explain how all this was resolved would take us far afield, so I’ll point to the references at the end of this section. When the dust settles, there is a well defined approximation that leads to Eq (890), and the resulting predictions can be made rigorous and shown to

be consistent with known bounds.

It would be good to connect these ideas with experiment. To what extent is our (or other animals’) performance in situations where we learn understandable in terms of these theoretical structures? A big problem here is what to measure. In the examples discussed above, what is being learned is a probability distribution, or some set of parameters describing the data that we observe. It’s not so easy to ask even a human subject to report on their current estimates of these parameters, and it’s completely unclear how we would do this in simpler organisms. In practice, subjects are usually asked to make a decision; in classical work on pigeons the decision is to peck or not to peck at a target, and for humans subjects are simply asked a yes/no question, or asked to push one of a small set of buttons. Evidently the bandwidth of these experiments is limited—although we may be continuously updating an internal model with many parameters, what we report is on the order of one bit, yes or no.

One context that comes closer to the theoretical discussion, albeit in a simple form, concerns making decisions when the alternatives come with unequal probabilities. This harkens back to our earliest topic, a human observer waiting for a dim flash of light in a dark room. As we discussed in that context, optimal decisions, deciding that a signal is convincingly above the background of noise, are achieved by setting a threshold that depends on the probability that the signal is present [\[need a definite pointer\]](#). If this probability can change over time, then it must be learned. More prosaically, if we have to choose between two alternatives even in a limit where they are fully distinguishable, but the rewards for the different choices vary probabilistically, then we have to learn something about the underlying probabilities of reward in order to develop a sensible strategy. These sorts of experiment have attracted interest because they might connect to our economic behavior, and because they provide settings in which we can search for the neural correlates of the subject’s estimate of probability and value.⁹⁰

There is a classical literature showing that human observers adjust their criteria for detecting signals to the probability that the signals occur. The question about learning is really how long it takes the subject to make this adjustment. In the simplest case, the probability changes suddenly, and we look for a change in behavior in response. If the only behavioral output is a decision among two alternatives, we as observers also need to go

⁹⁰ I think it is fair to say that the concept of “value” has attracted more attention in this context, because it seems more connected to economics. Indeed, there is now a whole field described as “neuro-economics.” But perhaps the probabilistic nature of our inferences, even in the economic context, have been given less attention than they deserve.

through an inference process to decide when is the first sign of a response. In such an experiment, we have a complete probabilistic description of the trajectory taken by the sensory stimuli or rewards, so at any moment we can calculate the probability that the signals being shown are consistent with constant parameters or a recent, sudden change. Given the responses of the subject, we can also ask for the moment at which we see the first statistical sign of a change in behavior. In experiments where rats experience changing reward probabilities, the change in behavior occurs at times so soon after the changes in probability that the best evidence for the change is modest, corresponding to probabilities in the range 0.1 to 0.9; only rarely (in $\sim 20\%$ of trials) do rats wait to reach 99% certainty. On these very short time scales, the rate at which the rat collects rewards changes very little, suggesting that changes in strategy really are driven by learning the underlying probabilities, rather than tinkering until rewards accumulate.

In a similar spirit, we can do a longer experiment, with the probabilities jumping among different levels, and track the dynamics of the behavior. [Have to decide how much to say here. Would like to connect with result from Corrado et al suggesting that filtering of experience to generate internal model of probability is near-optimal. Could do the calculation in a simple case, then point to




FIG. 164 Tracking the changing probabilities of reward, from Corrado et al (2005). At top, the local frequencies of choosing one of two alternatives (solid) and being rewarded (dashed), when the probability of this choice being rewarded jumps among different levels as shown (thin line); frequencies are computed from discrete events by smoothing with the Gaussian kernel shown in the inset. At bottom left, the filter inferred from the relationship between rewards and subsequent choices. At bottom right, the efficiency of collecting rewards averaged over the whole session, assuming that the subject implements a filter with the times constants as shown. The subjects' behaviors are best fit by parameters that generate efficiencies within one percent of the optimum.




FIG. 165 Learning curves in individuals vs groups (Gallistel et al 2004). [Need to give a full explanation of the experiment!] At left, average performance in a large population of birds improves gradually and very slowly, requiring many hundreds of trials before reaching its half maximal level. At right, performance measured in one individual bird is noisy (because we have access only to the number of pecks as a behavioral output), but makes a relatively sudden transition to near saturating performance as the animal experiences ~ 10 additional examples.

Fig 164 ... Need to digest the paper better, though.]

One approach to adding bandwidth to experiments on learning is to average over many subjects, so that the performance after N examples can be measured as a real number (e.g., the probability of subjects getting the right answer) even though the data from individuals is discrete (yes/no answers). But, as emphasized in Fig 165, this can be misleading. Individual subjects seem to learn simple tasks abruptly, but with transitions after different numbers of trials, so that average “learning curves” are smooth and gradual. This is interesting, because more abrupt learning reminds us of performance as a function of signal-to-noise ratio in discrimination tasks, and because theory along the lines described above often predicts relatively rapid learning when the space of possibilities is small. The variations across individuals may then reflect differences in how the ‘small problem’ posed by the particular experimental situation is weighted within the much larger set of possible behaviors available to the organism. But much needs to be done to make this precise.

Another approach to increasing the bandwidth of behavioral experiments is to look at continuous motor outputs rather than decisions. An example is if we have to move an object through a medium that generate an unknown, anisotropic mobility tensor; as we practice, we learn more about the parameters of our environment and can move more accurately. Importantly, each trial of such an experiment generates an entire movement trajectory

rather than just a single discrete decision. Analysis of these trajectories can reveal how the errors we make in one trial influence the change of our internal model on the next trial. [Should have a figure—maybe combine something from Shadmehr et al plus saccadic latency vs probability?] Although this emphasizes learning of parameters that influence the movement itself, the fact that some movements are made in extraordinary precise relations to sensory inputs (e.g., as we follow a moving target with our eyes), and that we can learn to anticipate the need for such movements (e.g., as targets follow predictable trajectories), suggests that analysis of continuous movements should more generally provide us with a path to examine more details of the brain’s internal model of the world. A simple version of this idea is that the latency for us to move our eyes toward one of two suddenly appearing targets depends on the relative probabilities of the targets—we move more quickly toward targets of higher probability, as shown in Fig 166, and it is tempting to think that the latency of movement gives us a readout of the brain’s estimate of this probability. Again, there is much to do here.

Thus far our examples of learning have been “passive.” That is, the learner experiences a data stream from which inferences can be drawn, but there is no way for the learner to shape the data stream, selecting observations which might be especially informative. [Give a discussion of infotaxis. This is interesting both as an active learning problem and as an example where gathering information substitutes effectively for “goal-directed behavior.”]

Finally, a theoretical point. We have emphasized that




FIG. 166 Latency for saccadic eye movements to targets of varying probability, from Carpenter & Williams (1995). Subjects are asked to move their eyes to a target which appears at a random time after they fixate on a small spot. The target is either to the left or right, with varying probabilities in blocks of trials. For two subjects (filled and empty symbols), one collects all the trials in which the subject move to a target of probability p , and computes the mean latency of the eye movement. Do we want to say anything about the distribution of latencies?

learning a model amounts to building an efficient representation of the data we have observed, and hence the “goal” of learning is no different than the goals proposed in the previous section for the transmission of information through neural or genetic networks. This theoretical unity is attractive. But one might worry—why do we care about representing what we have observed in the past? What matters, to follow the discussion at the end of Section IV.B, is what is of use in guiding our actions in the future. Thus, presumably we learn models that describe data collected in the past because we expect these models to still be true in the future, and this allows us to make successful predictions. How does this connect to our ideas about efficient representation?

We recall from Section IV.B that the predictive information in a time series, that is the information which observations on the past provide about the future, is equal to the subextensive component of the entropy. In the course of evaluating the probability of data given a class of models, in Eq (877), we have implicitly calculated this subextensive entropy. Specifically, we found that the negative log probability of a set of data at N time points had a term proportional to N (the extensive piece), and a term that grows only logarithmically with N (the leading subextensive piece), as in Eq [**]. Thus, when we are observing a time series from which we can learn a model with K parameters, there is a subextensive entropy and hence predictive information $\sim (K/2) \log_2 N$ bits. The “meaning” of this predictive information is precisely that we know something about the parameters underlying the data, and on the hypothesis that these parameters are constant we can predict something about the future.

Problem 168: Predictive information in learning. [One more problem with the details.]

When we observe N data points, the total amount of information we have collected is a number of bits proportional to N . But in the case we are considering, there are just $\sim (K/2) \log_2 N$ bits of information about the future. If we can separate these predictive bits from the nonpredictive background, we will have learned the parameters of the underlying model. Thus, compressing the data while preserving the predictive information is exactly the same problem as learning. Interestingly, if we live in a world described by a complex model (large K), then the amount of predictive information is much larger than the information needed to describe the present.

I think that our modern understanding of the preference for simple models, as explained here, is quite important, well known in certain circles, but less widely appreciated than it should be. Part of the difficulty is the presence of many independent threads in the literature. Rissanen had a very clear point of view which is essentially that presented here, although in different language; sources go back at least to Rissanen (1978), with a summary in Rissanen (1989). The problem became more urgent with the emergence of neural networks, which could be viewed as models with very large numbers of parameters. In this context, MacKay (1992) understood the critical role of ‘Occam factors,’ the integrals over parameter values that favor simpler models; see also his marvelous textbook (MacKay 2003). Balasubramanian (1997) generalized these ideas and translated them into physicists’ language, showing how the Occam factors can be thought of as entropy in the space of models. Certainly I learned a lot from talking to Balasubramanian, and from working out these ideas in the context of a field theoretic approach to learning distributions (Bialek et al 1996). For the case of the spherical horse, see Devine & Cohen (1992).

Balasubramanian 1997: Statistical inference, Occam’s razor, and statistical mechanics on the space of probability distributions. V Balasubramanian, *Neural Comp* **9**, 349–368 (1997).

Bialek et al 1996: Field theories for learning probability distributions. W Bialek, CG Callan & SP Strong, *Phys Rev Lett* **77**, 4693–4697 (1996).

Devine & Cohen 1992: *Absolute Zero Gravity: Science Jokes, Quotes and Anecdotes*. B Devine & JE Cohen (Fireside Press, Philadelphia, 1992).

MacKay 1992: A practical Bayesian framework for backpropagation networks. DJC MacKay, *Neural Comp* **4**, 448–472 (1992).

MacKay 2003: *Information Theory, Inference, and Learning Algorithms*. DJC MacKay (Cambridge University Press, Cambridge, 2003).

Rissanen 1978: Modeling by shortest data description. J Rissanen, *Automatica* **14**, 465–471 (1978).

Rissanen 1989: *Stochastic Complexity and Statistical Inquiry* J Rissanen (World Scientific, Singapore, 1989).

The study of neural networks also led to a very explicit formulation of learning as a statistical mechanics problem (Levin et al 1990). Within this framework it was discovered that there could be phase transitions in the learning of large models (Seung et al 1992), and these can be understood as a competition between energy (goodness of the fit) and entropy in the space of models; the effective temperature is the inverse of the number of examples we have seen, so the system ‘cools’ as we collect more data. Meanwhile the computer scientists have developed approaches to learning rules and distributions that focused on rigorous bounds—given that we have seen N examples, can we guarantee that our inferences are within ϵ of the correct model with probability $1 - \delta$? These ideas have their origins in Vapnik and Chernovonenkis (1971) and Valiant (1984). The rapprochement between the different approaches was given by Haussler et al (1996). For an early discussion about sudden vs. gradual learning, see Spence (1938). For a modern example, emphasizing the need for a unified approach to sudden and gradual learning, see Rubin et al (1997).

Haussler et al 1996: Rigorous learning curve bounds from statistical mechanics. D Haussler, M Kearns, HS Seung & N Tishby, *Machine Learning* **25**, 195–236 (1996).

Levin et al 1990: A statistical approach to learning and generalization in layered neural networks. E Levin, N Tishby & SA Solla, *Proc IEEE* **78**, 1568–1574 (1990).

Rubin et al 1997: Abrupt learning and retinal size specificity in illusory-contour perception. N Rubin, K Nakayama & R Shapley, *Curr Biol* **7**, 461–467 (1997).

Seung et al 1992: Statistical mechanics of learning from examples. HS Seung, H Sompolinsky & N Tishby, *Phys Rev A* **45**, 6056–6091 (1992).

Spence 1938: Gradual versus sudden solution of discrimination problems by chimpanzees. KW Spence, *J Comp Psych* **25**, 213–224 (1938).

Valiant 1984: A theory of the learnable. LG Valiant, *Commun ACM* **27**, 1134–1142 (1984).

Vapnik & Chernvonenkis 1971: On the uniform convergence of relative frequencies of events to their probabilities. VN Vapnik & AY Chervonenkis, *Theory of Probability and its Applications* **16**, 264–280 (1971).

Do we need a pointer back to Green and Swets, or some other reference about changing criteria in relation to changing probabilities?

The analysis of the response to sudden changes in probability is by Gallistel et al (2001). The experiments on fluctuating probabilities with primate subjects are by Sugrue et al (2004), and the analysis in Fig 164 is by Corrado et al (2005). For views on the emerging ideas of neuro-economics, see Glimcher (2003) and Camerer et al (2005). Measurements on learning through trial-by-trial analysis of continuous movement trajectories were pioneered by Thoroughman & Shadmehr (2000), who considered human arm movements; they had a particular view of the class of models that subjects use in these experiments, which simplified their analysis, but the idea is much more general. For measurements on the precision of tracking eye movements, see Osborne et al (2005, 2007). For the beautiful relationship between latency and target probability in saccadic eye movements, see Carpenter & Williams (1995). The idea of “infotaxis” is due to Vergassola et al (2007). Regarding the connection between learning and predictive information, see Bialek et al (2001) in Section IV.B.

Camerer et al 2005: Neuroeconomics: How neuroscience can inform economics. C Camerer, G Lowenstein & D Prelec, *J Econ Lit* **43**, 9–64 (2005).

Carpenter & Williams 1995: Neural computation of log likelihood in control of saccadic eye movements. RHS Carpenter & MLL Williams, *Nature* **377**, 59–62 (1995).

Corrado et al 2005: Linear–nonlinear–Poisson models of primate choice dynamics. GS Corrado, LP Sugrue, HS Seung & WT Newsome, *J Exp Anal Behav* **84**, 581–617 (2005).

Gallistel et al 2001: The rat approximates an ideal detector of changes in rates of reward: Implications for the law of effect. CR Gallistel, TA Mark, AP King & P Latham, *J Exp Psych: Animal Behav Proc* **27**, 354–372 (2001).

Gallistel et al 2004: The learning curve: Implications of a quantitative analysis. CR Gallistel, S Fairhurst & P Balsam, *Proc Nat’l Acad Sci (USA)* **101**, 13124–13131 (2004).

Glimcher 2003: *Decisions Uncertainty and the Brain: The Science of Neuroeconomics* PW Glimcher (MIT Press, Cambridge, 2003).

Osborne et al 2005: A sensory source for motor variation. LC Osborne, SG Lisberger & W Bialek, *Nature* **437**, 412–416 (2005).

Osborne et al 2007: Time course of precision in smooth pursuit eye movements of monkeys. LC Osborne, SS Hohl, W Bialek & SG Lisberger, *J Neurosci* **27**, 2987–2998 (2007).

Sugrue et al 2004: Matching behavior and the representation of value in the parietal cortex. LP Sugrue, GS Corrado & WT Newsome, *Science* **304**, 1782–1787 (2004).

Thoroughman & Shadmehr 2000: Learning of action through adaptive combination of motor primitives. KA Thoroughman & R Shadmehr, *Nature* **407**, 742–747 (2000).

Vergassola et al 2007: ‘Infotaxis’ as a strategy for searching without gradients. M Vergassola, E Villermaux & BI Shraiman, *Nature* **445**, 406–409 (2007).

E. Perspectives

Optimizing information transmission, or maximizing the efficiency with which information is represented, is the sort of abstract, general principle that physicists find appealing. At the same time, this abstraction makes us suspicious about its relevance to the nitty-gritty of life. Thus, while information is essential for survival, surely much of what organisms do is bound up in the fact that some bits are more useful than others, and in the challenges of acting rather than just collecting data. In this Chapter I have tried to show both how interesting predictions flow from the abstract principles, and how these principles connect, sometimes surprisingly, to the more quotidian facts of life. It is surely too early, in this as in any other section of the course, to decide if some candidate theoretical principles are “right,” and in any case I am not a disinterested observer. What I would like to emphasize here is that thinking about the optimization of information transmission has been productive, not least because it suggests genuinely new kinds of experiments. In many systems, these experiments have generated interesting results, independent of the theoretical motivation. In many other systems, even the first generation of experiments remains to be done.

Perhaps the most important point about information theoretic optimization principles is that they force us to think about biological systems in context. Whereas classical biology routinely considered organisms in their natural setting, as biology has modernized and become more mechanistic, we see more and more work on systems shorn of their context. To give an example, it may be that the best studied example of the regulation of gene expression is the *lac* operon in *E. coli*. But how much do we know about the distribution of lactose concentrations encountered by these cells in their natural environments? We know that, under many conditions, the total number of *lac* repressor proteins in the cell is small, but what is the dynamic range of this number over the lifetime of the organism? Vastly more is known about the details of the DNA sequences that are targeted by transcription factors involved in the regulation of metabolic genes than is known about the real world variations in nutrient conditions that create the need for metabolic regulation.

In the case of neural information processing, the ethologists—who often study systems specialized for the processing of particular sense data, such as bird song or bat echolocation—provided a persistent reminder about the importance of the natural context in understanding biological function. Perhaps our human abilities to deal with a seemingly much wider range of data and tasks generated some resistance to thinking that lessons from

a barn owl or an electric fish could be of relevance to how we explore higher brain function. The claim that at least some aspects of neural circuitry are arranged to generate efficient representations of incoming sense data provided a counterpoint, suggesting that even for a “general purpose” sensory system, context matters. By now there is a whole subfield of neuroscience focused on the structure and processing of natural signals, a field which we might think of as a modern, quantitative development of the early work in ethology. Because our sense organs are such high quality devices, there are substantial experimental challenges in characterizing their natural inputs and in delivering controlled versions of these natural signals in the laboratory. Precisely because natural signals are rich and complex, analyzing neural responses to these signals poses significant theoretical challenges (see, for example, Appendix A.7). Progress on these experimental and theoretical problems is giving us more powerful tools with which to explore the brain, again independent of the sometimes distant motivation from optimization principles.

Thinking about information flow encourages us to ask about the structure of natural behavioral outputs, as well as natural sensory inputs. In the attempt to quantify animal (and human) behavior in the laboratory, there has been a tradition of constraining this behavior to a small, discrete set of alternatives, and this has been enormously powerful, not least because such constrained experiments are amenable to analyses in terms of signals and noise as in our initial discussion of photon counting in vision. Similarly, experiments on the control of gene expression in single celled organisms often have focused on the “switch” in expression patterns associated with a sudden transition from one nutrient source to another. Even the ethologists tended to categorize, collapsing whole ranges of behavior onto a limited space of discrete choices. But behavior, from single cells to entire humans, is vastly richer than choosing among discrete alternatives. As the technology for monitoring behavior improves, it becomes possible to ask if the continuous variations in natural behaviors are just noise, or are related systematically to the goals and context. Even if behavior really is composed of choice among a small set of stereotyped possibilities—such as running and tumbling in *E. coli*—the timing of these choices can convey information about the sensory inputs that drive them.

We have the impression that we are bombarded by complex data, and that our behaviors are relatively limited. But the inputs to our sensory system are highly structured, presumably because they derive from a limited set of causes and effects in the environment, and hence carry much less information about what is really “out there” than one might guess from the available bandwidth; our receptors provide limited, noisy views of these inputs, reducing the information still further (see, for example, Problem [info in cone array]). At the other

end, our motor outputs in fact are quite rich, even if we tend to coarse grain and categorize these behaviors into limited classes. Could it be that motor outputs are so carefully shaped and timed in relation to sensory inputs from the environment that we (and other organisms) are making use of a large fraction of the information available about this environment? There is a huge experimental challenge in tracking information flow all the way from sensory input to motor output, even in simple cases, and in more complex cases there is a substantial theoretical challenge in providing a framework for the analysis of such data.

One of the most important aspects of information theory is the fact that bits have value. This is why, for example, there is a minimum number of bits we need to send over a telephone connection to be sure that speech is intelligible and speakers identifiable. For living organisms, the value of bits depends on many details, perhaps more detail than, as physicists, we would like to think about. What we can say, however, is that bits which have no predictive power are valueless, and that most of the bits we have collected over our lifetime are in this valueless category. Thus, separating predictive information from the background of non-predictive clutter is a formidable, and biologically relevant, challenge. Importantly, this very general task seems to contain within it, as special cases, problems ranging from signal processing to learning, problems that we usually think of as belonging to different levels of biological organization with very different mechanisms. Perhaps this is, after all, a path to the sort of general principle we are seeking.

V. OUTLOOK

This remains to be written. Evidently it requires a little distance from the details of the text, which I just haven't achieved yet! In broad terms, I want to summarize what we have seen, and point toward where we might be going. Return to some of Helmholtz' soaring ambition

Appendix A: Appendix

In these sections I collect things which are off the to side, or in the background, of the main arguments made in the text. Some readers will find the background essential, others will find the asides more interesting than what I thought were the main points. I hope, however, that everyone finds something useful here. As with the main text, I try not to skip steps, and problems are embedded in the narrative. **To a large extent, the Appendices are unedited as of September 18, 2011; for some of the newer ones, especially, much work is needed.**

1. Poisson processes

Photons from a conventional light source arrive at a detector as a random process, specifically a Poisson process. The defining feature of the Poisson process is that each event (photon arrival) is independent of all the others, given that we know the rate $r(t)$ at which the events occur. In these notes we'll go through the detailed consequences of this simple assumption of independence; hopefully some of the results are familiar. Note that many textbook presentations make a big deal out of the distinction between a “homogeneous” Poisson process, in which the rate is a constant, $r(t) = \bar{r}$, and an “inhomogeneous” Poisson process in which it can depend on time. The general case isn't that hard, so I prefer to start there.

One should perhaps note at the outset that most light sources are not exactly Poisson, but the approximation is very good. There are many more systems for which the Poisson model is a decent if not excellent approximation, and so we'll discuss all this without further reference to photons: we are describing the statistics of arbitrary point events which occur at times t_1, t_2, \dots, t_N .

The rate $r(t)$ can be thought of either as the mean rate of events that we would observe in the neighborhood of time t if we did the same experiment many times,

or equivalently as the probability per unit time that we observe an event at t . Recall that there is the same dual definition for the concentration $c(\mathbf{x})$ of molecules—either the mean number of molecule per unit volume that we find in the neighborhood of a point \mathbf{x} , or the probability per unit volume that we observe a single molecule at \mathbf{x} .

Since the events are independent, the probability density for observing events at times t_1, t_2, \dots, t_N must be proportional to a product of the rates evaluated at these times,

$$P[\{t_i\}|r(\tau)] \propto r(t_1)r(t_2)\cdots r(t_N) \equiv \prod_{i=1}^N r(t_i). \quad (\text{A1})$$

But to get the exact form of the distribution we must include a factor that measures the probability of *no* events occurring at any other times. The probability of an event occurring in a small bin of size $\Delta\tau$ surrounding time t is, by the original definition of the rate, $p(t) = r(t)\Delta\tau$, so the probability of no event must be $1 - p(t)$. Thus we need to form a product of factors $1 - p(t)$ for all times not equal to the special t_i where we observed events. Let's call this factor F ,

$$F = \prod_{n \neq i} [1 - p(t_n)]. \quad (\text{A2})$$

Then the probability of observing events in bins surrounding the t_i is

$$P[\{t_i\}|r(\tau)](\Delta\tau)^N = \frac{1}{N!} F \prod_{i=1}^N [r(t_i)\Delta\tau], \quad (\text{A3})$$

where the $N!$ corrects for all the different ways of assigning labels $1, 2, \dots, N$ to the events we observe.

To proceed we pull out all the factors related to the t_i and isolate the terms independent of these times:

$$\begin{aligned} P[\{t_i\}|r(\tau)](\Delta\tau)^N &= \frac{1}{N!} F \prod_{i=1}^N [r(t_i)\Delta\tau] \\ &= \frac{1}{N!} \prod_{n \neq i} [1 - r(t_n)\Delta\tau] \prod_{i=1}^N [r(t_i)\Delta\tau] \end{aligned} \quad (\text{A4})$$

$$= \frac{1}{N!} \prod_n [1 - r(t_n)\Delta\tau] \prod_{i=1}^N \left[\frac{r(t_i)\Delta\tau}{1 - r(t_i)\Delta\tau} \right]; \quad (\text{A5})$$

keep in mind that \prod_n denotes a product over *all* possible times t_n .

To simplify Eq (A5) we remember that products can be turned into sums by taking logarithms, so that

$$\prod_n [1 - r(t_n)\Delta\tau] = \exp \left(\sum_n \ln [1 - r(t_n)\Delta\tau] \right). \quad (\text{A6})$$

Now when we substitute back into Eq (A5) we find

$$\begin{aligned}
 P[\{t_i\}|r(\tau)](\Delta\tau)^N &= \frac{1}{N!} \prod_n [1 - r(t_n)\Delta\tau] \prod_{i=1}^N \left[\frac{r(t_i)\Delta\tau}{1 - r(t_i)\Delta\tau} \right] \\
 &= \frac{1}{N!} \exp \left(\sum_n \ln [1 - r(t_n)\Delta\tau] \right) \prod_{i=1}^N \left[\frac{r(t_i)\Delta\tau}{1 - r(t_i)\Delta\tau} \right]. \tag{A7}
 \end{aligned}$$

We are interested in the case where the time bin $\Delta\tau$ is very small (we introduced these artificially, remember), which means that we need to take the logarithm of numbers that are almost equal to one. We recall that the Taylor series of the logarithm is

$$\ln(1 + x) = x - \frac{1}{2}x^2 + \frac{1}{3}x^3 - \dots \tag{A8}$$

In this case we apply this expansion to

$$\ln [1 - r(t_n)\Delta\tau] = -r(t_n)\Delta\tau - \frac{1}{2} [r(t_n)\Delta\tau]^2 + \dots, \tag{A9}$$

so our expression for the probability can be written as

$$\begin{aligned}
 P[\{t_i\}|r(\tau)](\Delta\tau)^N &= \frac{1}{N!} \exp \left(\sum_n \ln [1 - r(t_n)\Delta\tau] \right) \prod_{i=1}^N \left[\frac{r(t_i)\Delta\tau}{1 - r(t_i)\Delta\tau} \right] \\
 &= \frac{1}{N!} \exp \left(\sum_n [-r(t_n)\Delta\tau] - \frac{1}{2} \sum_n [-r(t_n)\Delta\tau]^2 + \dots \right) \prod_{i=1}^N \left[\frac{r(t_i)\Delta\tau}{1 - r(t_i)\Delta\tau} \right]. \tag{A10}
 \end{aligned}$$

This expression involves a sum over bins, with factors of the bin width $\Delta\tau$. We recall that this converges, as the bins become small, to an integral:

$$\lim_{\Delta\tau \rightarrow 0} \sum_n f(t_n)\Delta\tau = \int dt f(t), \tag{A11}$$

for any smooth function $f(t)$. In the present case this means that

$$\lim_{\Delta\tau \rightarrow 0} \exp \left(\sum_n [-r(t_n)\Delta\tau] - \frac{1}{2} \sum_n [-r(t_n)\Delta\tau]^2 + \dots \right) = \exp \left[- \int dt r(t) - \frac{1}{2} \Delta\tau \int dt r^2(t) + \dots \right]. \tag{A12}$$

Now we notice that the second integral in the exponential has an extra factor of $\Delta\tau$, which comes from the $(\Delta\tau)^2$ in the previous expression, but if we really let $\Delta\tau$ go to zero this must be negligible as long as the rate doesn't become infinite.

Similarly, we have in Eq (A10) factors like

$$\frac{r(t_i)\Delta\tau}{1 - r(t_i)\Delta\tau},$$

and again as $\Delta\tau \rightarrow 0$ we can expand this in powers of $\Delta\tau$ and drop all but the first term. This is equivalent to replacing the denominator of the fraction by 1. So, when the dust clears, the expression for the probability density

of the event times becomes

$$P[\{t_i\}|r(\tau)] = \frac{1}{N!} \exp \left[- \int_0^T dt r(t) \right] \prod_{i=1}^N r(t_i), \tag{A13}$$

where we have set the limits on the integral to refer to the whole duration of our observations, from $t = 0$ to $t = T$. Note that this is a probability density for the N arrival times t_1, t_2, \dots, t_N and hence has units (time)^{-N}.

It is a useful exercise to check the normalization of the probability distribution in Eq. (A13). We want to calculate the total probability, which involves taking the term with N events and integrating over all N arrival times, then summing on N . Let's call this sum Z ,

$$Z \equiv \sum_{N=0}^{\infty} \int_0^T dt_1 \int_0^T dt_2 \cdots \int_0^T dt_N P[\{t_i\}|r(t)] \tag{A14}$$

$$= \sum_{N=0}^{\infty} \int_0^T dt_1 \int_0^T dt_2 \cdots \int_0^T dt_N \frac{1}{N!} \exp \left[- \int_0^T dt r(t) \right] \prod_{i=1}^N r(t_i). \tag{A15}$$

Notice that the exponential does not depend on the $\{t_i\}$ or on N , so we can take it outside the sum and integral. Furthermore, although we have to integrate over all the N different t_i together (an N dimensional integral), the

integrand is just a product of terms that depend on each individual t_i . This means that really we have a product of N one dimensional integrals:

$$Z = \exp \left[- \int_0^T dt r(t) \right] \sum_{N=0}^{\infty} \frac{1}{N!} \int_0^T dt_1 \cdots \int_0^T dt_N r(t_1) \cdots r(t_N) \tag{A16}$$

$$= \exp \left[- \int_0^T dt r(t) \right] \sum_{N=0}^{\infty} \frac{1}{N!} \int_0^T dt_1 r(t_1) \int_0^T dt_2 r(t_2) \cdots \int_0^T dt_N r(t_N) \tag{A17}$$

$$= \exp \left[- \int_0^T dt r(t) \right] \sum_{N=0}^{\infty} \frac{1}{N!} \left[\int_0^T dt r(t) \right]^N. \tag{A18}$$

Recall that the series expansion of the exponential function is

$$\exp(x) = \sum_{N=0}^{\infty} \frac{1}{N!} x^N, \tag{A19}$$

so we can actually do the sum in Eq. (A18):

$$\exp \left[- \int_0^T dt r(t) \right] \sum_{N=0}^{\infty} \frac{1}{N!} \left[\int_0^T dt r(t) \right]^N = \exp \left[- \int_0^T dt r(t) \right] \times \exp \left[+ \int_0^T dt r(t) \right] \tag{A20}$$

$$= 1, \tag{A21}$$

which completes our check on the normalization of the distribution.

Next we would like to derive an expression for the distribution of counts, which we write as $P(N|\langle N \rangle)$ to remind us that the shape of the distribution depends (as we will see) only on its mean. To do this we take the full probability distribution $P[\{t_i\}|r(\tau)]$, pick out the term involving N events, and then integrate over all the possible arrival times of these events:

$$P(N|\langle N \rangle) = \int_0^T dt_1 \cdots \int_0^T dt_N P[\{t_i\}|r(\tau)] \tag{A22}$$

$$= \int_0^T dt_1 \cdots \int_0^T dt_N \frac{1}{N!} \exp \left[- \int_0^T dt r(t) \right] \prod_{i=1}^N r(t_i). \tag{A23}$$

As in the discussion leading to Eq. (A18) we notice that the exponential factor can be taken outside the integral, and

that we have a product of N one dimensional integrals rather than a full N dimensional integral:

$$\begin{aligned} P(N|\langle N \rangle) &= \int_0^T dt_1 \cdots \int_0^T dt_N \frac{1}{N!} \exp \left[- \int_0^T dt r(t) \right] \prod_{i=1}^N r(t_i) \\ &= \frac{1}{N!} \exp \left[- \int_0^T dt r(t) \right] \int_0^T dt_1 \cdots \int_0^T dt_N \prod_{i=1}^N r(t_i) \\ &= \frac{1}{N!} \exp \left[- \int_0^T dt r(t) \right] \left[\int_0^T dt r(t) \right]^N \end{aligned} \quad (\text{A24})$$

$$\equiv \frac{1}{N!} \exp(-Q) Q^N, \quad (\text{A25})$$

where we have defined

$$Q = \int_0^T dt r(t). \quad (\text{A26})$$

In particular, the probability that no events occur in the time from $t = 0$ to $t = T$ is $P(0) = \exp(-Q)$, or

$$P(0|\langle N \rangle) = \exp \left[- \int_0^T dt r(t) \right]. \quad (\text{A27})$$

With the probability distribution of counts from Eq. (A25), we can compute the mean and the variance of the count. To obtain the mean we compute

$$\langle N \rangle \equiv \sum_{N=0}^{\infty} P(N) N \quad (\text{A28})$$

$$= \sum_{N=0}^{\infty} \frac{1}{N!} \exp(-Q) Q^N N \quad (\text{A29})$$

$$= \exp(-Q) \sum_{N=0}^{\infty} \frac{1}{N!} Q^N N. \quad (\text{A30})$$

Now we have already made use of the series expansion for the exponential, Eq. (A19), and to sum this last series we notice that

$$Q^N N = Q \frac{\partial}{\partial Q} Q^N, \quad (\text{A31})$$

so that

$$\begin{aligned} \langle N \rangle &= \exp(-Q) \sum_{N=0}^{\infty} \frac{1}{N!} Q^N N \\ &= \exp(-Q) \sum_{N=0}^{\infty} \frac{1}{N!} Q \frac{\partial}{\partial Q} Q^N \end{aligned} \quad (\text{A32})$$

$$= \exp(-Q) Q \frac{\partial}{\partial Q} \sum_{N=0}^{\infty} \frac{1}{N!} Q^N \quad (\text{A33})$$

$$= \exp(-Q) Q \frac{\partial}{\partial Q} \exp(+Q), \quad (\text{A34})$$

where in the last step we recognize the series for the exponential. Now the derivative of the exponential is just the exponential itself,

$$\frac{\partial}{\partial Q} \exp(+Q) = \exp(+Q), \quad (\text{A35})$$

so that

$$\begin{aligned} \langle N \rangle &= \exp(-Q) Q \frac{\partial}{\partial Q} \exp(+Q) \\ &= \exp(-Q) Q \exp(+Q) = Q. \end{aligned} \quad (\text{A36})$$

We see that the mean count is what we have called Q , the integral of the rate.

Now we can write the count distribution directly in terms of its mean:

$$P(N|\langle N \rangle) = \exp(-\langle N \rangle) \frac{\langle N \rangle^N}{N!}, \quad (\text{A37})$$

which is what we need to start the discussion of photon counting in vision, Eq (??).

We can do a very similar calculation to find the variance of the count distribution. We start by computing the average of N^2 ,

$$\langle N^2 \rangle = \sum_{N=0}^{\infty} N^2 P(N). \quad (\text{A38})$$

Substituting for $P(N)$ from Eq. (A25) and rearranging, we have

$$\begin{aligned} \langle N^2 \rangle &= \sum_{N=0}^{\infty} N^2 P(N) \\ &= \sum_{N=0}^{\infty} N^2 \exp(-Q) \frac{1}{N!} Q^N \end{aligned} \quad (\text{A39})$$

$$= \exp(-Q) \sum_{N=0}^{\infty} \frac{1}{N!} N^2 Q^N. \quad (\text{A40})$$

The trick is once again to write the extra factors of N (here N^2) in terms of derivatives with respect to Q . Now we know that

$$\frac{\partial^2}{\partial Q^2} Q^N = N(N-1)Q^{N-2}, \quad (\text{A41})$$

so we can write

$$Q^2 \frac{\partial^2}{\partial Q^2} Q^N = (N^2 - N)Q^N, \tag{A42}$$

which is almost what we want. But we can use the formula in Eq. (A31) to finish the job, obtaining

$$N^2 Q^N = Q^2 \frac{\partial^2}{\partial Q^2} Q^N + Q \frac{\partial}{\partial Q} Q^N. \tag{A43}$$

Now we can substitute into Eq. (A40) and follow the steps corresponding to Eq's (A32) through (A36):

$$\begin{aligned} \langle N^2 \rangle &= \exp(-Q) \sum_{N=0}^{\infty} \frac{1}{N!} N^2 Q^N \\ &= \exp(-Q) \sum_{N=0}^{\infty} \frac{1}{N!} \left[Q^2 \frac{\partial^2}{\partial Q^2} Q^N + Q \frac{\partial}{\partial Q} Q^N \right] \end{aligned} \tag{A44}$$

$$= \exp(-Q) Q^2 \frac{\partial^2}{\partial Q^2} \sum_{N=0}^{\infty} \frac{1}{N!} Q^N + \exp(-Q) Q \frac{\partial}{\partial Q} \sum_{N=0}^{\infty} \frac{1}{N!} Q^N \tag{A45}$$

$$= \exp(-Q) Q^2 \frac{\partial^2}{\partial Q^2} \exp(+Q) + \exp(-Q) Q \frac{\partial}{\partial Q} \exp(+Q) \tag{A46}$$

$$= \exp(-Q) Q^2 \exp(+Q) + \exp(-Q) Q \exp(+Q) \tag{A47}$$

$$= Q^2 + Q. \tag{A48}$$

Now since we have already identified Q as equal to the mean count, this means that the mean square count can be written as

$$\langle N^2 \rangle = \langle N \rangle^2 + \langle N \rangle. \tag{A49}$$

But the variance of the count is defined by

$$\langle (\delta N)^2 \rangle \equiv \langle N^2 \rangle - \langle N \rangle^2 \tag{A50}$$

$$= [\langle N \rangle^2 + \langle N \rangle] - \langle N \rangle^2 = \langle N \rangle. \tag{A51}$$

Thus the variance of the count for a Poisson process is equal to the mean count.

The next characteristic of the Poisson process is the interval between events. The probability per unit time that we observe an event at time t is given by the rate, $r(t)$. The probability that we observe no events in the interval $[t, t + \tau)$ is given by

$$P(0) = \exp \left[- \int_t^{t+\tau} dt' r(t') \right]. \tag{A52}$$

The probability per unit time that this interval is closed by an event is again the rate, now at time $t + \tau$. Thus the probability per unit time that we see events at t and $t + \tau$, with no events in between is given by

$$P(t, t + \tau) = r(t) \exp \left[- \int_t^{t+\tau} dt' r(t') \right] r(t + \tau). \tag{A53}$$

In the simple case that the rate is constant, this is just $P(t, t + \tau) = r^2 e^{-r\tau}$. On the other hand, if the rate

varies, the average probability for observing two events separated by an empty interval of duration τ is

$$P_2(\tau) = \left\langle r(t) \exp \left[- \int_t^{t+\tau} dt' r(t') \right] r(t + \tau) \right\rangle, \tag{A54}$$

where $\langle \dots \rangle$ is an average over these variations in rate.

If we ask for the probability density of intervals, this is really the conditional probability that the next event will be at $t + \tau$ given that there was an event at t . To form this conditional probability we need to divide by the probability of an event at t , but this is just the average rate. Again, in the simple case of constant rate, this yields the probability density of inter-event intervals,

$$p(\tau) = r e^{-r\tau}. \tag{A55}$$

This exponential form is one of the classic signatures of a Poisson process. We can think of it as arising because the moment at which the interval closes has no memory of the moment at which it opened, and so the probability that there has not been an event must be a product of terms for the absence of an event in each small time slice $\Delta\tau$, as in the derivation above, and this product becomes an exponential.

Our last task is to evaluate averages over Poisson processes, such as the one in Eq (33),

$$\left\langle \sum_i V_0(t - t_i) \right\rangle = \sum_{N=0}^{\infty} \int_0^T dt_1 \cdots \int_0^T dt_N P[\{t_i\}|r(t)] \sum_i V_0(t - t_i). \quad (\text{A56})$$

We proceed simply and systematically, looking at one term in our sum and doing the integrals one at a time.

One term in the sum means that we choose, for ex-

ample $i = 1$ and one particular value of N . This term is

$$\int_0^T dt_1 \cdots \int_0^T dt_N P[\{t_i\}|r(t)] V_0(t - t_1) = \int_0^T dt_1 \cdots \int_0^T dt_N \exp \left[- \int_0^T d\tau r(\tau) \right] \frac{1}{N!} r(t_1) r(t_2) \cdots r(t_N) V_0(t - t_1). \quad (\text{A57})$$

Notice that the exponential factor (along the the $1/N!$) is constant and comes outside the integral. Now we rearrange the order of the integrals:

$$\int_0^T dt_1 \int_0^T dt_2 \cdots \int_0^T dt_N r(t_1) r(t_2) \cdots r(t_N) V_0(t - t_1) = \int_0^T dt_1 r(t_1) V_0(t - t_1) \int_0^T dt_2 r(t_2) \cdots \int_0^T dt_N r(t_N) \quad (\text{A58})$$

$$= \left[\int_0^T dt_1 r(t_1) V_0(t - t_1) \right] \left[\int_0^T d\tau r(\tau) \right]^{N-1}. \quad (\text{A59})$$

But the fact that we chose $i = 1$ was arbitrary; we would have gotten the same answer for any $i = 1, 2, \dots, N$. Thus summing over i is the same as multiplying by N . This leaves us with the sum on N , so we put everything back together to find

$$\left\langle \sum_i V_0(t - t_i) \right\rangle = \exp \left[- \int_0^T d\tau r(\tau) \right] \left[\int_0^T dt_1 r(t_1) V_0(t - t_1) \right] \sum_{N=0}^{\infty} \frac{N}{N!} \left[\int_0^T d\tau r(\tau) \right]^{N-1} \quad (\text{A60})$$

$$= \exp \left[- \int_0^T d\tau r(\tau) \right] \int_0^T dt_1 r(t_1) V_0(t - t_1) \sum_{N=0}^{\infty} \frac{1}{N!} \left[\int_0^T d\tau r(\tau) \right]^N \quad (\text{A61})$$

$$= \exp \left[- \int_0^T d\tau r(\tau) \right] \int_0^T dt_1 r(t_1) V_0(t - t_1) \exp \left[+ \int_0^T d\tau r(\tau) \right] \quad (\text{A62})$$

$$= \int_0^T dt_1 r(t_1) V_0(t - t_1). \quad (\text{A63})$$

Thus what we have shown is that our simple model of summing pulses from single photons generates a voltage that responds linearly to the light intensity,

$$\langle V(t) \rangle = V_{\text{DC}} + \int dt' V_0(t - t') r(t'), \quad (\text{A64})$$

which is Eq (34) in the main text.

Actually, we have shown something more general, which will be useful below. The expectation value we have computed is of the form

$$\left\langle \sum_i f(t_i) \right\rangle. \quad (\text{A65})$$

What we have seen is that summing over arrival times is,

on average, equivalent to integrating over the rate,

$$\left\langle \sum_i f(t_i) \right\rangle = \int d\tau r(\tau) f(\tau). \quad (\text{A66})$$

Intuitively, this makes sense: the sum over arrival times approximates a density along the time axis, and this density is the rate, with units of (events)/(time).

Now we need to do the same calculation, but for the correlation function of the voltage. Again we have

$$V(t) = \sum_i V_0(t - t_i), \quad (\text{A67})$$

and we want to compute $\langle V(t)V(t') \rangle$. Intuitively, the arrival times of photons are independent of one another—

this is the essence of the Poisson process—and so we should have

$$\langle V(t)V(t') \rangle = \left\langle \sum_i V_0(t-t_i) \sum_j V_0(t'-t_j) \right\rangle \quad (\text{A68})$$

$$= \sum_{ij} \langle V_0(t-t_i)V_0(t'-t_j) \rangle \quad (\text{A69})$$

$$= \sum_{i \neq j} \langle V_0(t-t_i)V_0(t'-t_j) \rangle + \sum_i \langle V_0(t-t_i)V_0(t'-t_i) \rangle \quad (\text{A70})$$

$$= \sum_{i \neq j} \langle V_0(t-t_i) \rangle \langle V_0(t'-t_j) \rangle + \sum_i \langle V_0(t-t_i)V_0(t'-t_i) \rangle, \quad (\text{A71})$$

where we use the independence of t_i and t_j for $i \neq j$ in the last step. It's useful to add and subtract the “diagonal” $i = j$ term from the sum, so that

$$\begin{aligned} \langle V(t)V(t') \rangle &= \sum_{i \neq j} \langle V_0(t-t_i) \rangle \langle V_0(t'-t_j) \rangle + \sum_i \langle V_0(t-t_i) \rangle \langle V_0(t'-t_i) \rangle \\ &\quad + \sum_i \langle V_0(t-t_i)V_0(t'-t_i) \rangle - \sum_i \langle V_0(t-t_i) \rangle \langle V_0(t'-t_i) \rangle \end{aligned} \quad (\text{A72})$$

$$= \sum_{ij} \langle V_0(t-t_i) \rangle \langle V_0(t'-t_j) \rangle + \sum_i [\langle V_0(t-t_i)V_0(t'-t_i) \rangle - \langle V_0(t-t_i) \rangle \langle V_0(t'-t_i) \rangle]. \quad (\text{A73})$$

The key step now is to notice that we can rearrange the sums and expectation values in the first term,

$$\sum_{ij} \langle V_0(t-t_i) \rangle \langle V_0(t'-t_j) \rangle = \sum_i \langle V_0(t-t_i) \rangle \sum_j \langle V_0(t'-t_j) \rangle \quad (\text{A74})$$

$$= \left\langle \sum_i V_0(t-t_i) \right\rangle \left\langle \sum_j V_0(t'-t_j) \right\rangle \quad (\text{A75})$$

$$= \langle V(t) \rangle \langle V(t') \rangle, \quad (\text{A76})$$

where in the last step we recognize the voltage itself, from Eq (A67). Thus Eq (A73) can be rewritten as an equation for the covariance of the voltage fluctuations,

$$\langle \delta V(t)\delta V(t') \rangle \equiv \langle V(t)V(t') \rangle - \langle V(t) \rangle \langle V(t') \rangle = \sum_i [\langle V_0(t-t_i)V_0(t'-t_i) \rangle - \langle V_0(t-t_i) \rangle \langle V_0(t'-t_i) \rangle]. \quad (\text{A77})$$

If we confine our attention to the simple case where the rate is constant, $r(t) = \bar{r}$, then the second term in brackets must be a constant, since $\langle V_0(t-t_i) \rangle$ involves averaging over all possible times t_i , and with constant rate all these times are equally likely. So, if we don't worry about constants, we can write

$$\langle \delta V(t)\delta V(t') \rangle \sim \sum_i \langle V_0(t-t_i)V_0(t'-t_i) \rangle \quad (\text{A78})$$

$$= \left\langle \sum_i V_0(t-t_i)V_0(t'-t_i) \right\rangle, \quad (\text{A79})$$

and now we can use Eq (A66) to give

$$\langle \delta V(t)\delta V(t') \rangle = \bar{r} \int d\tau V_0(t-\tau)V_0(t'-\tau), \quad (\text{A80})$$

where again we are neglecting a constant.

It is especially useful to convert the correlation function of voltage fluctuations into the corresponding power spectrum, since then any uncertainties about constants will go away. More precisely, if we had a constant term in the correlation function it would show up as a term $\sim \delta(\omega)$ in the power spectrum, and all we need to do is to be sure that we drop any such terms. In general, the power spectrum is

$$S_V(\omega) = \int d\tau e^{+i\omega\tau} \langle \delta V(t+\tau)\delta V(t) \rangle, \quad (\text{A81})$$

and so in this case we have

$$S_V(\omega) = \int d\tau e^{+i\omega\tau} \bar{r} \int d\tau' V_0(t + \tau - \tau') V_0(t - \tau') \quad (\text{A82})$$

$$= \bar{r} \int d\tau \int d\tau' e^{+i\omega\tau} e^{+i\omega(t-\tau')} V_0(t + \tau - \tau') e^{-i\omega(t-\tau')} V_0(t - \tau') \quad (\text{A83})$$

$$= \bar{r} \left[\int d\tau e^{+i\omega(\tau+t-\tau')} V_0(\tau + t - \tau') \right] \left[\int d\tau' e^{-i\omega(t-\tau')} V_0(t - \tau') \right] \quad (\text{A84})$$

$$= \bar{r} \left| \tilde{V}_0(\omega) \right|^2, \quad (\text{A85})$$

where in the last step we recognize the Fourier transform of the pulse shape $V_0(t)$. This is what we need for Eq (58) of the main text.

Problem 169: More carefully. Fill in the details of the calculation above, being sure to keep track of the floating constants. Verify that, when you are careful, there is no term $\sim \delta(\omega)$ in the power spectrum. Can you generalize this discussion to the case of time varying rates?

Portions of this section were adapted from Rieke et al (1997). The connection between power spectra and the shape of single photon (or more general Poisson) events is sometimes called Campell's theorem, and there is a classic discussion by Rice (1944–45), reprinted in the marvelous book edited by Wax (1954); the other articles in this book (by Chandrasekar and others) also are very much worth reading! Feynman & Hibbs (1965) give a beautiful discussion of how a Poisson stream of pulses comes to approximate continuous, Gaussian noise; of course there is much more in this book as well. For a more complete discussion of photon statistics, and the role of coherent states, one can look to yet another classic paper, Glauber (1963).

Feynman & Hibbs 1965: *Quantum Mechanics and Path Integrals*. RP Feynman & AR Hibbs (McGraw-Hill, New York, 1965).

Glauber 1963: Coherent and incoherent states of the radiation field. RJ Glauber, *Phys Rev* **131**, 2766–2788 (1963).

Rice 1944–45: Mathematical analysis of random noise. SO Rice, *Bell Sys Tech J* **23**, 282–332 (1944) & **24**, 46–156 (1945).

Rieke et al 1997: *Spikes: Exploring the Neural Code*. F Rieke, D Warland, RR de Ruyter van Steveninck & W Bialek (MIT Press, Cambridge, 1997).

Wax 1954: *Selected Papers on Noise and Stochastic Processes*. N Wax, ed (Dover Publications, New York, 1954).

2. Correlations, power spectra and all that

Consider a function $x(t)$ that varies in time. We would like to describe a situation in which these variations are random, drawn out of some distribution. But now we need a distribution for a function, rather than for a finite set of variables. This shouldn't bother us, since such constructions are central to much of modern physics, for example in the path integral approach to quantum mechanics. We refer to distributions of functions as “distribution functionals” when we need to be precise.

One strategy for constructing distribution functionals is to start by discretizing time, so that we have at most a countable infinity of variables $x(t_1), x(t_2), x(t_3), \dots$. Let's assume for simplicity that the mean value of x is zero. Then the first nontrivial characterization of the statistics of x is the covariance matrix,

$$C_{ij} = \langle x(t_i)x(t_j) \rangle. \quad (\text{A86})$$

We recall that if a single variable y is drawn from a Gaussian distribution with zero mean, then we have

$$P(y) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{y^2}{2\sigma^2}\right]. \quad (\text{A87})$$

The generalization to multiple variables is

$$P(\{x_i\}) = \frac{1}{\sqrt{(2\pi)^N \det C}} \exp\left[-\frac{1}{2} \sum_{i,j=1}^N x_i (C^{-1})_{ij} x_j\right], \quad (\text{A88})$$

where as usual \det is the determinant and $(C^{-1})_{ij}$ is the ij element of the matrix inverse to C ; if we think of the $\{x_i\}$ as a vector \mathbf{x} , then we can write, more compactly,

$$P(\{x_i\}) = \frac{1}{\sqrt{(2\pi)^N \det C}} \exp\left[-\frac{1}{2} \mathbf{x}^T \cdot C^{-1} \cdot \mathbf{x}\right], \quad (\text{A89})$$

where \mathbf{x}^T is the transpose of the vector \mathbf{x} . Just to be clear, this describes a Gaussian distribution, but we have no guarantee that \mathbf{x} will be Gaussian.

Problem 170: Gaussian integrals. If you haven't done these before, now is a good time to check that the probability distribution




FIG. 167 Covariance matrix and its inverse. At left, the covariance matrix in Eq (A94), with $\Delta t/\tau_c = 0.1$. At right, the inverse matrix, with inset showing a 10×10 submatrix surrounding the diagonal.

in Eq (A89) is normalized. This requires you to show that

$$\int d^N x \exp \left[-\frac{1}{2} \mathbf{x}^T \cdot C^{-1} \cdot \mathbf{x} \right] = \sqrt{(2\pi)^N \det C}. \quad (\text{A90})$$

While you're at it, you should also show that

$$\ln \det C = \text{Tr} \ln C. \quad (\text{A91})$$

This should be straightforward for the case which matters here, where C must have well defined, positive eigenvalues.

In general the covariance matrix C_{ij} can have an arbitrary structure, constrained only by symmetry and positivity of its eigenvalues. But when the index i refers to discrete time points, we have an extra constraint that comes from invariance under translations in time. Because there is no clock, we must have that

$$\langle x(t)x(t') \rangle = C_x(t - t'), \quad (\text{A92})$$

with no dependence on the absolute time t or t' . As an example, if

$$C_x(t - t') = e^{-|t-t'|/\tau_c}, \quad (\text{A93})$$

and $t_n = n\Delta t$, then

$$C_{ij} = \exp \left[-\left(\frac{\Delta t}{\tau_c} \right) |i - j| \right]. \quad (\text{A94})$$

This is shown in Fig 167 for $\Delta t/\tau_c = 0.1$.

It is useful to look directly at the inverse matrix, also shown in Fig 167. We see that this inverse matrix consists almost entirely of zeros, except in the immediate neighborhood of the diagonal. This tells us that the inverse matrix actually is the discretization of a differential operator. Reflexively, seeing that we have to compute inverses and determinants of matrices, we should think about diagonalizing C . We recall from quantum mechanics that the eigenfunctions of an operator have to provide a representation of the underlying symmetries. In this case, the relevant symmetry is time translation, so we know to look at the Fourier functions, $e^{-i\omega t}$. In fact, once we have the hint that we should use a Fourier representation, we don't need the crutch of discrete time points any more. Let's see how this works.

We define the Fourier transform with the conventions

$$\tilde{x}(\omega) = \int_{-\infty}^{\infty} dt e^{+i\omega t} x(t), \quad (\text{A95})$$

$$x(t) = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} e^{-\omega t} \tilde{x}(\omega). \quad (\text{A96})$$

Now if we compute the covariance of two frequency components, we have

$$\langle \tilde{x}(\omega)\tilde{x}(\omega') \rangle = \left\langle \int_{-\infty}^{\infty} dt e^{+i\omega t} x(t) \int_{-\infty}^{\infty} dt' e^{+i\omega' t'} x(t') \right\rangle \quad (\text{A97})$$

$$= \int_{-\infty}^{\infty} dt e^{+i\omega t} \int_{-\infty}^{\infty} dt' e^{+i\omega' t'} \langle x(t)x(t') \rangle \quad (\text{A98})$$

$$= \int_{-\infty}^{\infty} dt e^{+i\omega t} \int_{-\infty}^{\infty} dt' e^{+i\omega' t'} \int_{-\infty}^{\infty} \frac{d\Omega}{2\pi} e^{-i\Omega(t-t')} S_x(\Omega), \quad (\text{A99})$$

where we introduce the Fourier transform of the correlation function,

$$S_x(\Omega) = \int_{-\infty}^{\infty} d\tau e^{+i\Omega\tau} C_x(\tau) \quad (\text{A100})$$

$$C_x(t - t') = \int_{-\infty}^{\infty} \frac{d\Omega}{2\pi} e^{-i\Omega(t-t')} S_x(\Omega). \quad (\text{A101})$$

Now we can rearrange the integrals in Eq (A99):

$$\begin{aligned}\langle \tilde{x}(\omega)\tilde{x}(\omega') \rangle &= \int_{-\infty}^{\infty} dt e^{+i\omega t} \int_{-\infty}^{\infty} dt' e^{+i\omega' t'} \int_{-\infty}^{\infty} \frac{d\Omega}{2\pi} e^{-i\Omega(t-t')} S_x(\Omega), \\ &= \int_{-\infty}^{\infty} \frac{d\Omega}{2\pi} S_x(\Omega) \left[\int_{-\infty}^{\infty} dt e^{i(\omega-\Omega)t} \right] \left[\int_{-\infty}^{\infty} dt' e^{i(\omega'+\Omega)t'} \right].\end{aligned}\tag{A102}$$

This is moment to recall the Fourier representation of the Dirac delta function. The delta function has the property that

$$\delta(z) = 0 \quad z \neq 0, \tag{A103}$$

$$\int dz \delta(z) = 1, \tag{A104}$$

if the domain of the integral includes $z = 0$. Then

$$\delta(z) = \int_{-\infty}^{\infty} \frac{dq}{2\pi} e^{-iqz}. \tag{A105}$$

Thus we recognize, in Eq (A102),

$$\int_{-\infty}^{\infty} dt e^{i(\omega-\Omega)t} = 2\pi\delta(\omega - \Omega), \tag{A106}$$

$$\int_{-\infty}^{\infty} dt' e^{i(\omega'+\Omega)t'} = 2\pi\delta(\omega' + \Omega). \tag{A107}$$

Substituting back into Eq (A102), we have

$$\begin{aligned}\langle \tilde{x}(\omega)\tilde{x}(\omega') \rangle &= \int_{-\infty}^{\infty} \frac{d\Omega}{2\pi} S_x(\Omega) \left[\int_{-\infty}^{\infty} dt e^{i(\omega-\Omega)t} \right] \left[\int_{-\infty}^{\infty} dt' e^{i(\omega'+\Omega)t'} \right]. \\ &= \int_{-\infty}^{\infty} \frac{d\Omega}{2\pi} S_x(\Omega) 2\pi\delta(\omega - \Omega) 2\pi\delta(\omega' + \Omega) \\ &= S_x(\omega) 2\pi\delta(\omega' + \omega).\end{aligned}\tag{A108}$$

We see that, while different time points can be correlated with one another in complicated ways, the covariance of frequency components has a much simpler structure: $\tilde{x}(\omega)$ is correlated only with $\tilde{x}(-\omega)$.

This covariance structure, which couples positive and negative frequency components, makes sense when we realize that we are using a complex representation for real variables. To make a real variable $x(t)$, the Fourier transform must obey

$$\tilde{x}(-\omega) = \tilde{x}^*(\omega), \tag{A110}$$

so positive and negative frequency components are not independent—in fact they are redundant. It might be more natural to write Eq (A109) as

$$\langle \tilde{x}(\omega)\tilde{x}^*(\omega') \rangle = S_x(\omega) 2\pi\delta(\omega' - \omega), \tag{A111}$$

making clear that frequency components are correlated with themselves, not with other frequencies.

We could instead think about the real and imaginary parts of the positive frequency components, which can be written as

$$\tilde{x}_{\text{Re}}(\omega) = \frac{1}{2} [\tilde{x}(\omega) + \tilde{x}(-\omega)], \tag{A112}$$

and

$$\tilde{x}_{\text{Im}}(\omega) = \frac{1}{2i} [\tilde{x}(\omega) - \tilde{x}(-\omega)]. \tag{A113}$$

With this representation, we can use the result in Eq (A109):

$$\langle \tilde{x}_{\text{Re}}(\omega)\tilde{x}_{\text{Re}}(\omega') \rangle = \left\langle \frac{1}{2} [\tilde{x}(\omega) + \tilde{x}(-\omega)] \frac{1}{2} [\tilde{x}(\omega') + \tilde{x}(-\omega')] \right\rangle \tag{A114}$$

$$= \frac{1}{4} [\langle \tilde{x}(\omega)\tilde{x}(\omega') \rangle + \langle \tilde{x}(\omega)\tilde{x}(-\omega') \rangle + \langle \tilde{x}(-\omega)\tilde{x}(\omega') \rangle + \langle \tilde{x}(-\omega)\tilde{x}(-\omega') \rangle] \tag{A115}$$

$$= \frac{S_x(\omega)}{4} 2\pi [\delta(\omega + \omega') + \delta(\omega - \omega') + \delta(-\omega + \omega') + \delta(-\omega - \omega')]. \tag{A116}$$

Because we are looking only at positive frequencies, $\omega + \omega'$ can never be zero, and hence the first and last delta functions can be dropped. The remaining two are actually the same, so we have

$$\langle \tilde{x}_{\text{Re}}(\omega) \tilde{x}_{\text{Re}}(\omega') \rangle = \frac{1}{2} S_x(\omega) 2\pi \delta(\omega - \omega'). \quad (\text{A117})$$

Similar calculations show that the imaginary parts of $\tilde{x}(\omega)$ have the same variance,

$$\begin{aligned} \langle \tilde{x}_{\text{Im}}(\omega) \tilde{x}_{\text{Im}}(\omega') \rangle &= \langle \tilde{x}_{\text{Re}}(\omega) \tilde{x}_{\text{Re}}(\omega') \rangle & (\text{A118}) \\ &= \frac{1}{2} S_x(\omega) 2\pi \delta(\omega - \omega'), & (\text{A119}) \end{aligned}$$

while real and imaginary parts are uncorrelated,

$$\langle \tilde{x}_{\text{Re}}(\omega) \tilde{x}_{\text{Im}}(\omega') \rangle = 0. \quad (\text{A120})$$

Problem 171: The other phase. Derive Eq's (A119) and (A120).

$$P[\{\tilde{x}_{\text{Re}}(\omega)\}] \propto \exp \left[-\frac{1}{2} \int_0^\infty \frac{d\omega}{2\pi} \int_0^\infty \frac{d\omega'}{2\pi} \tilde{x}_{\text{Re}}(\omega) \mathcal{A}(\omega, \omega') \tilde{x}_{\text{Re}}(\omega') \right], \quad (\text{A121})$$

where \mathcal{A} is the inverse of the covariance,

$$\int \frac{d\omega'}{2\pi} \mathcal{A}(\omega, \omega') \langle \tilde{x}_{\text{Re}}(\omega') \tilde{x}_{\text{Re}}(\omega'') \rangle = 2\pi \delta(\omega - \omega''). \quad (\text{A122})$$

We can find \mathcal{A} by substituting the explicit expression for the covariance and doing the integrals:

$$\begin{aligned} 2\pi \delta(\omega - \omega'') &= \int \frac{d\omega'}{2\pi} \mathcal{A}(\omega, \omega') \langle \tilde{x}_{\text{Re}}(\omega') \tilde{x}_{\text{Re}}(\omega'') \rangle \\ &= \int \frac{d\omega'}{2\pi} \mathcal{A}(\omega, \omega') \frac{1}{2} S_x(\omega') 2\pi \delta(\omega' - \omega'') & (\text{A123}) \end{aligned}$$

$$= \frac{1}{2} \mathcal{A}(\omega, \omega'') S_x(\omega''). \quad (\text{A124})$$

What does all this mean? We think of the random function of time $x(t)$ as being built out of frequency components, and each component has a real and imaginary part. The structure of the covariance matrix is such that different frequency components do not covary, and this makes sense—if we have covariation of different frequency components then we can beat them against each other to make a clock running at the difference frequency, and this would violate time translation invariance. Similarly, the fact that real and imaginary components do not covary means that there is no preferred phase, which again is consistent with (indeed, required by) time translation invariance.

We should be able to put these results on the covariance matrix together to describe the distribution functional for a Gaussian function of time. Since the real and imaginary parts are independent, let's start with just the real parts. We should have

Thus, we have

$$\mathcal{A}(\omega, \omega'') = \frac{1}{S_x(\omega'')} 4\pi \delta(\omega - \omega''). \quad (\text{A125})$$

Substituting back into Eq (A121) for the probability distribution, we have

$$P[\{\tilde{x}_{\text{Re}}(\omega)\}] \propto \exp \left[-\frac{1}{2} \int_0^\infty \frac{d\omega}{2\pi} \int_0^\infty \frac{d\omega'}{2\pi} \tilde{x}_{\text{Re}}(\omega) \mathcal{A}(\omega, \omega') \tilde{x}_{\text{Re}}(\omega') \right], \quad (\text{A126})$$

$$= \exp \left[-\frac{1}{2} \int_0^\infty \frac{d\omega}{2\pi} \int_0^\infty \frac{d\omega'}{2\pi} \tilde{x}_{\text{Re}}(\omega) \frac{4\pi \delta(\omega - \omega'')}{S_x(\omega'')} \tilde{x}_{\text{Re}}(\omega') \right] \quad (\text{A127})$$

$$= \exp \left[-\int_0^\infty \frac{d\omega}{2\pi} \frac{\tilde{x}_{\text{Re}}^2(\omega)}{S_x(\omega)} \right]. \quad (\text{A128})$$

Exactly the same argument applies to the imaginary

parts of the Fourier components, and these are indepen-

dent of the real parts, so we have

$$P[x(t)] = P[\{\tilde{x}_{\text{Re}}(\omega), \tilde{x}_{\text{Im}}(\omega)\}] \quad (\text{A129})$$

$$\propto \exp \left[- \int_0^\infty \frac{d\omega}{2\pi} \frac{\tilde{x}_{\text{Re}}^2(\omega) + \tilde{x}_{\text{Im}}^2(\omega)}{S_x(\omega)} \right] \quad (\text{A130})$$

$$= \frac{1}{Z} \exp \left[- \int_0^\infty \frac{d\omega}{2\pi} \frac{|\tilde{x}(\omega)|^2}{S_x(\omega)} \right] \quad (\text{A131})$$

$$= \frac{1}{Z} \exp \left[- \frac{1}{2} \int_{-\infty}^\infty \frac{d\omega}{2\pi} \frac{|\tilde{x}(\omega)|^2}{S_x(\omega)} \right], \quad (\text{A132})$$

where we have introduced the normalization constant Z .

It's useful to look at the example illustrated in Fig 167. Here we have $C_x(\tau) = \exp(-|\tau|/\tau_c)$, so the power spectrum is

$$S_x(\omega) = \int_{-\infty}^\infty d\tau e^{i\omega\tau} e^{-|\tau|/\tau_c} \quad (\text{A133})$$

$$= \int_{-\infty}^0 d\tau e^{(+i\omega+1/\tau_c)\tau} + \int_0^\infty d\tau e^{(+i\omega-1/\tau_c)\tau} \quad (\text{A134})$$

$$= \frac{1}{(+i\omega+1/\tau_c)} + \frac{1}{-(+i\omega-1/\tau_c)} \quad (\text{A135})$$

$$= \frac{\tau_c}{1+i\omega\tau_c} + \frac{\tau_c}{1-i\omega\tau_c} \quad (\text{A136})$$

$$= \frac{2\tau_c}{1+(\omega\tau_c)^2}. \quad (\text{A137})$$

This means that the probability distribution functional has the form

$$\begin{aligned} P[x(t)] &= \frac{1}{Z} \exp \left[- \frac{1}{2} \int_{-\infty}^\infty \frac{d\omega}{2\pi} \frac{|\tilde{x}(\omega)|^2}{S_x(\omega)} \right] \\ &= \frac{1}{Z} \exp \left[- \frac{1}{4\tau_c} \int_{-\infty}^\infty \frac{d\omega}{2\pi} [1 + (\omega\tau_c)^2] |\tilde{x}(\omega)|^2 \right]. \end{aligned} \quad (\text{A138})$$

We recall that

$$\int_{-\infty}^\infty \frac{d\omega}{2\pi} |\tilde{x}(\omega)|^2 = \int dt x^2(t). \quad (\text{A139})$$

More subtly,

$$\int_{-\infty}^\infty \frac{d\omega}{2\pi} (\omega\tau_c)^2 |\tilde{x}(\omega)|^2 = \tau_c^2 \int_{-\infty}^\infty \frac{d\omega}{2\pi} | -i\omega\tilde{x}(\omega) |^2 \quad (\text{A140})$$

$$= \tau_c^2 \int dt \left[\frac{dx(t)}{dt} \right]^2, \quad (\text{A141})$$

where we recognize $-i\omega\tilde{x}(\omega)$ as the Fourier transform of $dx(t)/dt$. Thus we can write

$$\begin{aligned} P[x(t)] &= \frac{1}{Z} \exp \left[- \frac{1}{4\tau_c} \int_{-\infty}^\infty \frac{d\omega}{2\pi} [1 + (\omega\tau_c)^2] |\tilde{x}(\omega)|^2 \right] \\ &= \frac{1}{Z} \exp \left[- \frac{1}{4\tau_c} \int dt (\tau_c^2 \dot{x}^2(t) + x^2(t)) \right] \end{aligned} \quad (\text{A142})$$

This shows explicitly, as promised above, that inverting the covariance matrix gives rise to differential operators. This example also is nice because it produces a probability distribution functional for trajectories $x(t)$ that reminds us of a (Euclidean) path integral in quantum mechanics, in this case for the harmonic oscillator.

Let's push a little further and see if we can evaluate the normalization constant Z . By definition, we have

$$Z = \int \mathcal{D}x \exp \left[- \frac{1}{4\tau_c} \int dt (\tau_c^2 \dot{x}^2(t) + x^2(t)) \right], \quad (\text{A143})$$

where $\int \mathcal{D}x$ denotes an integral over all the functions $x(t)$. We have the general result for an N dimensional Gaussian integral,

$$\int d^N x \exp \left[- \frac{1}{2} \mathbf{x}^T \cdot \hat{A} \cdot \mathbf{x} \right] = \sqrt{\frac{(2\pi)^N}{\det \hat{A}}} \quad (\text{A144})$$

$$= \sqrt{(2\pi)^N} \exp \left[- \frac{1}{2} \text{Tr} \ln \hat{A} \right], \quad (\text{A145})$$

where \hat{A} is a matrix. Here we need to let the number of dimensions become infinite, since we are integrating over functions. As you may recall from discussions of the path integral in quantum mechanics, there is some arbitrariness about how we do this, or, more formally, in how we define the measure $\mathcal{D}x$. A fairly standard choice is to absorb the $\sqrt{2\pi}$, so that, in the time window $0 < t < T$,

$$\mathcal{D}x = \lim_{dt \rightarrow 0} \prod_{n=0}^{T/dt} \frac{dx(t_n)}{\sqrt{2\pi}}, \quad t_n = n \cdot dt. \quad (\text{A146})$$

Notice that before we send $dt \rightarrow 0$, we have an integral over a finite number of points, so we should be able to carry over the results we know, and just interpret the limits correctly.

The Gaussian functional integrals that we want to do have the general form

$$\int \mathcal{D}x \exp \left[- \frac{1}{2} \int dt \int dt' x(t) \hat{K}(t, t') x(t') \right],$$

where \hat{K} is an operator. Carrying over what we know from the case of finite matrices [Eq (A145)], we have

$$\int \mathcal{D}x \exp \left[-\frac{1}{2} \int dt \int dt' x(t) \hat{K}(t, t') x(t') \right] = \exp \left[-\frac{1}{2} \text{Tr} \ln \hat{K} \right]. \tag{A147}$$

Our only problem is to say what we mean by $\text{Tr} \ln \hat{K}$. Since \hat{K} is an operator, we can ask for its spectrum, that is the eigenvalues and eigenfunctions. This means that we need to solve the equations

$$\int_0^T dt' \hat{K}(t, t') u_\mu(t') = \Lambda_\mu u_\mu(t), \tag{A148}$$

where we are careful here to note that we are working in window $0 < t < T$. In the basis formed by the eigenfunctions, of course \hat{K} is diagonal. As with matrices, when an operator is diagonal we can take the log element by element, and then computing the trace requires us to sum over these diagonal elements; recall that traces and determinants are invariant, so we can use this convenient

basis and not worry about generality. Thus,

$$\text{Tr} \ln \hat{K} = \sum_\mu \ln \Lambda_\mu. \tag{A149}$$

How does this work for our case? First, we need to identify the operator \hat{K} . In the exponential of $P[x(t)]$ we have

$$\int dt \left[\frac{\tau_c}{2} \left(\frac{dx(t)}{dt} \right)^2 + \frac{1}{2\tau_c} x^2(t) \right].$$

To get this into a more standard form we need to integrate by parts,

$$\int dt \left[\frac{\tau_c}{2} \left(\frac{dx(t)}{dt} \right)^2 + \frac{1}{2\tau_c} x^2(t) \right] = \int dt x(t) \left[-\frac{\tau_c}{2} \frac{d^2}{dt^2} + \frac{1}{2\tau_c} \right] x(t). \tag{A150}$$

This allows us to identify We now see that our integral for Z in Eq (A143) can be written

$$\hat{K}(t', t) = \delta(t' - t) \left[-\frac{\tau_c}{2} \frac{d^2}{dt^2} + \frac{1}{2\tau_c} \right]. \tag{A151}$$

This is a linear operator, and also time translation invariant (again). So we know that the eigenfunctions are $e^{-i\omega t}$, and since we are in a finite window of duration T we should use only those frequency components that ‘fit’ into the window, $\omega_n = 2\pi n/T$ for integer n . We have

$$\int_0^T dt \delta(t' - t) \left[-\frac{\tau_c}{2} \frac{d^2}{dt^2} + \frac{1}{2\tau_c} \right] e^{-i\omega_n t} = \left(\frac{\tau_c \omega_n^2}{2} + \frac{1}{2\tau_c} \right) e^{-i\omega_n t'}, \tag{A152}$$

so that the eigenvalues are

$$\Lambda(\omega_n) = \left(\frac{\tau_c \omega_n^2}{2} + \frac{1}{2\tau_c} \right) = \frac{1 + (\omega_n \tau_c)^2}{2\tau_c}. \tag{A153}$$

Notice that these are just the inverses of the power spectrum,

$$\Lambda(\omega_n) = \frac{1}{S_x(\omega_n)}. \tag{A154}$$

This makes sense, of course, when we look back at Eq (A132).

To finish the calculation, we have

$$Z = \exp \left[-\frac{1}{2} \sum_\mu \Lambda_\mu \right] \tag{A155}$$

$$= \exp \left[-\frac{1}{2} \sum_n \ln \left(\frac{1}{S_x(\omega_n)} \right) \right] \tag{A156}$$

$$= \exp \left[\frac{1}{2} \sum_n \ln S_x(\omega_n) \right]. \tag{A157}$$

Finally, we need to do the sum. As the time window T becomes large, the spacing between frequency components, $\Delta\omega = 2\pi/T$, become small, and we expect that the

sum approaches an integral.⁹¹ Thus, for any function of ω_n ,

$$\sum_n f(\omega_n) = \frac{1}{\Delta\omega} \sum_n \Delta\omega f(\omega_n) \quad (\text{A159})$$

$$\rightarrow \frac{1}{\Delta\omega} \int d\omega f(\omega) \quad (\text{A160})$$

$$= T \int \frac{d\omega}{2\pi} f(\omega). \quad (\text{A161})$$

At last, this gives us

$$Z = \exp \left[\frac{T}{2} \int \frac{d\omega}{2\pi} \ln S_x(\omega) \right]. \quad (\text{A162})$$

Putting the pieces together, we have the probability distribution functional for a Gaussian $x(t)$,

$$P[x(t)] = \exp \left[+\frac{T}{2} \int_{-\infty}^{\infty} \ln S_x(\omega) - \frac{1}{2} \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \frac{|\tilde{x}(\omega)|^2}{S_x(\omega)} \right]. \quad (\text{A163})$$

Not every case we look at will be Gaussian, but this helps to get us started.

Problem 172: Generality. We made an effort to evaluate Z in the specific case where $C_x(\tau) = e^{-|\tau|/\tau_c}$, but we wrote the final result in a very general form, Eq (A163). Show that this slide into generality was justified.

Problem 173: Nonzero means and signal to noise ratios. We should be able to carry everything through in the case where the mean $x(t)$ is not zero. For example, if we just have background noise described by some spectrum $\mathcal{N}(\omega)$, then

$$P_{\text{noise}}[x(t)] = \exp \left[+\frac{T}{2} \int_{-\infty}^{\infty} \ln \mathcal{N}(\omega) - \frac{1}{2} \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \frac{|\tilde{x}(\omega)|^2}{\mathcal{N}(\omega)} \right]. \quad (\text{A164})$$

If there is an added signal $x_0(t)$, the distribution functional becomes

$$P_{\text{signal}}[x(t)] = \exp \left[+\frac{T}{2} \int_{-\infty}^{\infty} \ln \mathcal{N}(\omega) - \frac{1}{2} \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \frac{|\tilde{x}(\omega) - \tilde{x}_0(\omega)|^2}{\mathcal{N}(\omega)} \right]. \quad (\text{A165})$$

Suppose that you observe some particular $x(t)$, and you have to decide whether this came from the signal or noise distribution, that is, you have to decide whether the signal was present; for simplicity assume that the two possibilities are equally likely a priori. As discussed in Chapter 1, to make such decisions optimally you should use the relative probabilities that the signal or noise could give rise to your data. In particular, consider computing the “log likelihood ratio,”

$$\lambda[x(t)] \equiv \ln \left(\frac{P_{\text{signal}}[x(t)]}{P_{\text{noise}}[x(t)]} \right) \quad (\text{A166})$$

⁹¹ There is an analogous result for summing over the states of particles in a box in quantum systems; recall that the states are labelled by their wavevector \mathbf{k} , and in three dimensions we have

$$\sum_{\mathbf{k}} \rightarrow V \int \frac{d^3k}{(2\pi)^3}, \quad (\text{A158})$$

where V is the volume of the box.

(a.) Give a simple expression for $\lambda[x(t)]$. Show that it is a linear functional of $x(t)$.

(b.) Show that, when the $x(t)$ are drawn at random out of either P_{signal} or P_{noise} , $\lambda[x(t)]$ is a Gaussian random variable. Find the means, $\langle \lambda \rangle_{\text{noise}}$ and $\langle \lambda \rangle_{\text{signal}}$, and the variances $\langle (\delta\lambda)^2 \rangle_{\text{noise}}$ and $\langle (\delta\lambda)^2 \rangle_{\text{signal}}$, in the two distributions. Hint: you should see that $\langle (\delta\lambda)^2 \rangle_{\text{noise}} = \langle (\delta\lambda)^2 \rangle_{\text{signal}}$.

(c.) Sketch the distributions $P_{\text{noise}}(\lambda)$ and $P_{\text{signal}}(\lambda)$. Show that your ability to make reliable discriminations is determined only by the signal to noise ratio,

$$SNR = \frac{\langle (\lambda)_{\text{signal}} - \langle \lambda \rangle_{\text{noise}} \rangle^2}{\langle (\delta\lambda)^2 \rangle}, \quad (\text{A167})$$

and that we can write

$$SNR = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \frac{|\tilde{x}_0(\omega)|^2}{\mathcal{N}(\omega)}. \quad (\text{A168})$$

(d.) In rod cells, a single photon produces a current pulse with the approximate form $x_0(t) = I_1(t/\tau)^3 e^{-t/\tau}$. The power spectrum of continuous background noise is approximately $\mathcal{N}(\omega) = A/[1 + (\omega\tau)^2]^2$, with the same value of τ . Evaluate the peak current, I_{peak} , and total variance of the background noise, σ_I^2 . A naive estimate of the signal to noise ratio is just $SNR_{\text{naive}} = (I_{\text{peak}}/\sigma_I)^2$. Show that the optimal signal to noise ratio, computed from Eq (A168), is larger. Why?

Is there anything more to say here? Maybe some discussion of “states” of molecules and correlation functions? Perhaps some references.

3. Electronic transition in large molecules

In this section we’ll outline an honest calculation that reproduces the intuition of Figs 19 and 20. We have a system with two electronic states, which we can represent as a spin one-half; let spin down be the ground state and spin up be the excited state. The Born–Oppenheimer approximation tells us that we can think of the atoms in a molecule as moving in a potential determined by the electronic state,⁹² which we denote by $V_{\uparrow}(\mathbf{q})$ and $V_{\downarrow}(\mathbf{q})$ in the excited and ground states, respectively; \mathbf{q} stands for all the atomic coordinates (not just the one in the sketches above). Since we are observing photon absorption, there must be a matrix element that connects the two electronic states and couples to the electromagnetic field; we’ll assume that, absent symmetries, this coupling is dominated by an electric dipole term. In principle the dipole matrix element \vec{d} could depend upon the atomic coordinates, but we’ll neglect this effect.⁹³ Putting the pieces together, we have the Hamiltonian for the molecule

$$\mathbf{H} = \mathbf{K} + \frac{1}{2}(1 + \sigma_z)V_{\uparrow}(\mathbf{q}) + \frac{1}{2}(1 - \sigma_z)V_{\downarrow}(\mathbf{q}) + \vec{d} \cdot \vec{E}(\sigma_+ + \sigma_-), \quad (\text{A169})$$

⁹² As in the main text, I’ll use “atoms” and “nuclei” interchangeably.

⁹³ In practice, this is a small effect. You should think about why this is true.

where \mathbf{K} is the kinetic energy of the atoms. To this we should of course add the usual Hamiltonian for the electromagnetic field.

We are interested in computing the rate at which photons of energy $\hbar\Omega$ are absorbed, and of course we will do this as a perturbation expansion in the term $\sim \vec{d}$. The result of such a calculation can be presented as the ‘Golden rule’ for transition rates, but this formulation hides the underlying dynamics. So, at the risk of being pedantic, I’ll go through the steps that usually lead to the Golden rule and take a detour that leads us to a formula in which the dynamics of atomic motions are more explicit.⁹⁴

We start our system in the ground state of the electrons ($|\downarrow\rangle$), in some initial state ($|i\rangle$) of the atomic coordinates, and in the presence of one photon of wavevector \vec{k} and frequency $\Omega = c|\vec{k}|$ (polarization is an unnecessary complication here). As the system evolves under the Hamiltonian \mathbf{H} , at some time t we want to measure the probability of finding the system in the excited state $|\uparrow\rangle$, in some other state of the atoms $|f\rangle$, and absent the photon. The general statement is that quantum states evolve as

$$|\psi(0)\rangle \rightarrow |\psi(t)\rangle = \mathbf{T} \exp \left[-\frac{i}{\hbar} \int_0^t d\tau \mathbf{H}(\tau) \right] |\psi(0)\rangle, \tag{A170}$$

where \mathbf{T} is the time ordering operator. Thus, for our particular problem, the probability of starting in state $|\downarrow, i, \vec{k}\rangle$ and ending in state $|\uparrow, f, \emptyset\rangle$ is given by

$$p_{i \rightarrow f}(t) = \left| \langle \emptyset, f, \uparrow | \mathbf{T} \exp \left[-\frac{i}{\hbar} \int_0^t d\tau \mathbf{H}(\tau) \right] | \downarrow, i, \vec{k} \rangle \right|^2. \tag{A171}$$

In fact, we don’t care about the final state of the atoms, and we can’t select their initial state—this comes out of the Boltzmann distribution. So we really should compute

$$P(t) = \sum_{i,f} \left| \langle \emptyset, f, \uparrow | \mathbf{T} \exp \left[-\frac{i}{\hbar} \int_0^t d\tau \mathbf{H}(\tau) \right] | \downarrow, i, \vec{k} \rangle \right|^2 p_i, \tag{A172}$$

where p_i is the probability of being in the initial atomic state i .

As usual, we will break the Hamiltonian into two pieces, $\mathbf{H} = \mathbf{H}_0 + \mathbf{H}_1$, and do perturbation theory in \mathbf{H}_1 . We choose $\mathbf{H}_1 = \vec{d} \cdot \vec{E}(\sigma_+ + \sigma_-)$, which is the only term that connects the states $|\downarrow\rangle$ and $|\uparrow\rangle$. The leading term in the perturbation series thus becomes

$$P(t) \approx \frac{1}{\hbar^2} \sum_{i,f} \left| \langle \emptyset, f, \uparrow | \mathbf{T} e^{-\frac{i}{\hbar} \int_0^t d\tau \mathbf{H}_0(\tau)} \int_0^t d\tau' \mathbf{H}_1(\tau') | \downarrow, i, \vec{k} \rangle \right|^2 p_i. \tag{A173}$$

If we look more carefully at the amplitude, we have

$$\langle \emptyset, f, \uparrow | \mathbf{T} e^{-\frac{i}{\hbar} \int_0^t d\tau \mathbf{H}_0(\tau)} \int_0^t d\tau' \mathbf{H}_1(\tau') | \downarrow, i, \vec{k} \rangle = \int_0^t d\tau' \langle f | \mathbf{T} \left(e^{-\frac{i}{\hbar} \int_{\tau'}^t d\tau \mathbf{H}_0(\tau)} \right) \vec{d} \cdot \langle \emptyset | \vec{E} | \vec{k} \rangle \mathbf{T} \left(e^{-\frac{i}{\hbar} \int_0^{\tau'} d\tau \mathbf{H}_0(\tau)} \right) | i \rangle e^{-i\Omega\tau'}, \tag{A174}$$

where τ' is the moment at which the term $\mathbf{H}_1 \sim \sigma_+$ acts to flip the state from $|\downarrow\rangle$ to $|\uparrow\rangle$; the terms $\mathbf{H}_{\downarrow,\uparrow}$ are defined by

$$\mathbf{H}_{\downarrow} = \mathbf{K} + V_{\downarrow}(\mathbf{q}), \tag{A175}$$

and similarly for \mathbf{H}_{\uparrow} . The key point is that when we square this amplitude and sum over final states, we can identify this as a sum over a complete set of states, and we recall that

$$\sum_f |f\rangle \langle f| = \mathbf{1}, \tag{A176}$$

the unit operator. Further, to keep things simple, let’s assume that motion of the atoms is approximately classical. Because the terms $\mathbf{H}_{\uparrow,\downarrow}$ depend only on the atomic coordinates and momenta, the classical approximation means that we don’t have to worry about the non-commutativity of these operators at different times, and we can drop the formalities of time ordering. Putting all of the terms together, we can rewrite $P(t)$ from Eq (A173):

$$P(t) \approx \frac{(\vec{d} \cdot \langle \emptyset | \vec{E} | \vec{k} \rangle)^2}{\hbar^2} \int_0^t d\tau_1 \int_0^t d\tau_2 e^{+i\Omega(\tau_1 - \tau_2)} \sum_i p_i \langle i | e^{+\frac{i}{\hbar} \int_0^{\tau_1} d\tau \mathbf{H}_\downarrow(\tau)} e^{+\frac{i}{\hbar} \int_{\tau_1}^t d\tau \mathbf{H}_\uparrow(\tau)} e^{-\frac{i}{\hbar} \int_{\tau_2}^t d\tau \mathbf{H}_\uparrow(\tau)} e^{-\frac{i}{\hbar} \int_0^{\tau_2} d\tau \mathbf{H}_\downarrow(\tau)} | i \rangle \quad (\text{A177})$$

$$= \frac{(\vec{d} \cdot \langle \emptyset | \vec{E} | \vec{k} \rangle)^2}{\hbar^2} \int_0^t d\tau_1 \int_0^t d\tau_2 e^{+i\Omega(\tau_1 - \tau_2)} \sum_i p_i \langle i | \exp \left(+\frac{i}{\hbar} \int_{\tau_1}^{\tau_2} d\tau [\mathbf{H}_\uparrow(\tau) - \mathbf{H}_\downarrow(\tau)] \right) | i \rangle \quad (\text{A178})$$

$$\propto \int_0^t d\tau_1 \int_0^t d\tau_2 e^{+i\Omega(\tau_1 - \tau_2)} \left\langle \exp \left[+\frac{i}{\hbar} \int_{\tau_1}^{\tau_2} d\tau \epsilon[\mathbf{q}(\tau)] \right] \right\rangle, \quad (\text{A179})$$

where $\epsilon = \mathbf{H}_\uparrow - \mathbf{H}_\downarrow = V_\uparrow - V_\downarrow$ is the instantaneous energy difference between the ground and excited states, which fluctuates as the atomic coordinates fluctuate, and $\langle \dots \rangle$ denotes an average over these fluctuations.

Problem 174: Missing steps. Fill in the steps leading to Eq (A179). If you are more ambitious, try the case where the atomic motions are fully quantum mechanical.

Notice that the integrand in Eq (A179) depends only on the time difference $\tau_2 - \tau_1$. Thus, we are doing an integral of the form

$$\int_0^t d\tau_1 \int_0^t d\tau_2 F(\tau_2 - \tau_1). \quad (\text{A180})$$

It seems natural to rewrite this integral over the (τ_1, τ_2) plane in terms of an integral over the time difference and the mean. In the limit that t is large, this yields

$$\int_0^t d\tau_1 \int_0^t d\tau_2 F(\tau_2 - \tau_1) \rightarrow t \int_{-\infty}^{\infty} d\tau F(\tau). \quad (\text{A181})$$

Thus, we have

$$P(t) \propto t \int_{-\infty}^{\infty} d\tau e^{+i\Omega\tau} \left\langle \exp \left[-\frac{i}{\hbar} \int_0^\tau d\tau' \epsilon[\mathbf{q}(\tau')] \right] \right\rangle, \quad (\text{A182})$$

so that the transition rate or absorption cross-section for photons of frequency Ω becomes

$$\sigma(\Omega) \propto \left\langle \int_{-\infty}^{\infty} d\tau \exp \left[+i\Omega\tau - \frac{i}{\hbar} \int_0^\tau d\tau' \epsilon[\mathbf{q}(\tau')] \right] \right\rangle. \quad (\text{A183})$$

Now we can recover the intuition of Fig 19 as a saddle point approximation to the integral in Eq (A183). We recall that the saddle point approximation is

$$\int dt \exp[+i\phi(t)] \approx \sqrt{\frac{2\pi}{|\phi''(t_*)|}} \exp[+i\phi(t_*)], \quad (\text{A184})$$

where the time t_* is defined by

$$\left. \frac{d\phi(t)}{dt} \right|_{t=t_*} = 0. \quad (\text{A185})$$

The condition for validity of the approximation is that the time scale

$$\delta t \sim 1/\sqrt{|\phi''(t_*)|} \quad (\text{A186})$$

be small compared with the intrinsic time scales for variation of $\phi(t)$. As applied to Eq (A183), the saddle point condition is

$$0 = \left. \frac{d}{d\tau} \left[+i\Omega\tau - \frac{i}{\hbar} \int_0^\tau d\tau' \epsilon[\mathbf{q}(\tau')] \right] \right|_{\tau=\tau_*} \quad (\text{A187})$$

$$= i\Omega - \frac{i}{\hbar} \epsilon[\mathbf{q}(\tau_*)] \quad (\text{A188})$$

$$\hbar\Omega = \epsilon[\mathbf{q}(\tau_*)]. \quad (\text{A189})$$

Thus, the saddle point condition states that the integral defining the cross-section is dominated by moments when the instantaneous difference between the ground and excited state energies matches the photon energy. But this instantaneous difference $\epsilon[\mathbf{q}]$ is exactly the ‘vertical’ energy difference in Fig 19. Since this integral is inside an expectation value over the fluctuations in atomic coordinates, the cross-section will be proportional to the probability that this matching condition is obeyed.

If the sketch in Fig 19 is equivalent to a saddle point approximation, we have to consider conditions for validity of this approximation. The time scale defined by Eq (A186) becomes

$$\delta t \sim \left| \frac{1}{\hbar} \frac{d\epsilon[\mathbf{q}(\tau)]}{d\tau} \right|^{-1/2} \sim \sqrt{\frac{\hbar}{\epsilon'v}}, \quad (\text{A190})$$

where ϵ' is the slope of the energy difference as a function of atomic coordinates, and v is a typical velocity for motion along these coordinates. Thus large slopes result in smaller values of δt , and of course this time scales as $\sqrt{\hbar}$. The natural time scale of motion along the atomic coordinates is given by vibrational periods,

or $\omega_{\text{vib}}^{-1} = \tau_{\text{vib}} \sim \Delta/v$, where Q is a typical displacement from equilibrium. This lets us write

$$\delta t \sim \sqrt{\frac{\hbar}{\epsilon'v}} \sim \sqrt{\frac{\hbar\omega_{\text{vib}}}{\epsilon'Q} \cdot \frac{Q}{v\omega_{\text{vib}}}} \sim \tau_{\text{vib}} \sqrt{\frac{\hbar\omega_{\text{vib}}}{\epsilon'Q}}. \quad (\text{A191})$$

We see that $\delta t \ll \tau_{\text{vib}}$ if the energy $\epsilon'Q$ is much larger than the energy of vibrational quanta $\hbar\omega_{\text{vib}}$. But $\epsilon'Q$ is the range of energy differences between the ground and excited states that the molecule can access as it fluctuates—and this is the width of the absorption spectrum. Thus, self-consistently, if we find that the width of the spectrum is large compared to the vibrational quanta, then our saddle point approximation is accurate.

We can go a bit further if we specialize to the case where, as in Fig 20, the different potential surfaces are exactly Hookean springs, that is when the dynamics of atomic motions are harmonic oscillators. In the general case there are many normal modes, so we would write

$$V_{\uparrow}(\mathbf{q}) = \frac{1}{2} \sum_i \omega_i^2 q_i^2 \quad (\text{A192})$$

$$V_{\downarrow}(\mathbf{q}) = \epsilon_0 + \frac{1}{2} \sum_i \omega_i^2 (q_i - \Delta_i)^2. \quad (\text{A193})$$

In this case,

$$\epsilon[\mathbf{q}(t)] \equiv V_{\uparrow}[\mathbf{q}(t)] - V_{\downarrow}[\mathbf{q}(t)] \quad (\text{A194})$$

$$= \epsilon_0 + \frac{1}{2} \sum_i \omega_i^2 \Delta_i^2 - \sum_i \omega_i^2 \Delta_i q_i(t) \quad (\text{A195})$$

$$= \hbar\Omega_{\text{peak}} - X(t), \quad (\text{A196})$$

where the generalized coordinate $X(t)$ is given by a weighted combination of all the modes,

$$X(t) = \sum_i \omega_i^2 \Delta_i q_i(t). \quad (\text{A197})$$

Equation (A183) for the absorption cross-section thus becomes

$$\begin{aligned} \sigma(\Omega) &\propto \left\langle \int_{-\infty}^{\infty} d\tau \exp \left[+i\Omega\tau - \frac{i}{\hbar} \int_0^{\tau} d\tau' \epsilon[\mathbf{q}(\tau')] \right] \right\rangle \\ &= \left\langle \int_{-\infty}^{\infty} d\tau \exp \left[+i\Omega\tau - \frac{i}{\hbar} \int_0^{\tau} d\tau' (\hbar\Omega_{\text{peak}} - X(\tau')) \right] \right\rangle \\ &= \int_{-\infty}^{\infty} d\tau e^{+i(\Omega - \Omega_{\text{peak}})\tau} \left\langle \exp \left[+\frac{i}{\hbar} \int_0^{\tau} d\tau' X(\tau') \right] \right\rangle. \end{aligned} \quad (\text{A198})$$

The key point is that, because $X(t)$ is a sum of harmonic oscillator coordinates, its fluctuations are drawn from a Gaussian distribution when we compute the average $\langle \dots \rangle$ over the equilibrium ensemble.

Problem 175: Gaussian averages. Derive Eq (A199).

We recall that, if y is a Gaussian random variable, then

$$\langle e^y \rangle = \exp \left[\langle y \rangle + \frac{1}{2} \langle (\delta y)^2 \rangle \right]. \quad (\text{A199})$$

In the present case, the role of y is played by an integral over the trajectory of $X(t)$, but this shouldn't bother us:

$$\begin{aligned} &\left\langle \exp \left[+\frac{i}{\hbar} \int_0^{\tau} d\tau' X(\tau') \right] \right\rangle \\ &= \exp \left[\frac{1}{2} \left\langle \left(\frac{i}{\hbar} \int_0^{\tau} d\tau' X(\tau') \right)^2 \right\rangle \right] \quad (\text{A200}) \\ &= \exp \left[-\frac{1}{2\hbar^2} \int_0^{\tau} d\tau_1 \int_0^{\tau} d\tau_2 \langle X(\tau_1) X(\tau_2) \rangle \right], \end{aligned} \quad (\text{A201})$$

where we start by making use of the fact that $\langle X \rangle = 0$.

We see from Eq (A201) that the shape of the absorption spectrum is determined by the correlation function of the modes to which the electronic transition are coupled, that is $C_X(\tau_1 - \tau_2) = \langle X(\tau_1) X(\tau_2) \rangle$. If these modes have relatively slow dynamics, then the time scales τ that enter the integral we need to do will be much shorter than

the time scales over which this correlation function varies. In this limit we can approximate

$$\int_0^\tau d\tau_1 \int_0^\tau d\tau_2 \langle X(\tau_1)X(\tau_2) \rangle \approx \int_0^\tau d\tau_1 \int_0^\tau d\tau_2 \langle X(0)X(0) \rangle = \langle X^2 \rangle \tau^2. \quad (\text{A202})$$

Notice also that

$$\langle X^2 \rangle = \left\langle \left(\sum_i \omega_i^2 \Delta_i q_i \right)^2 \right\rangle = \sum_i \omega_i^4 \Delta_i^2 \langle q_i^2 \rangle; \quad (\text{A203})$$

in the classical limit we have $\langle q_i^2 \rangle = k_B T / \omega_i^2$, and hence

$$\langle X^2 \rangle = k_B T \sum_i \omega_i^2 \Delta_i^2 = 2k_B T \lambda, \quad (\text{A204})$$

where λ generalizes the reorganization energy or Stokes' shift to the case of many modes. Finally, putting these pieces together, we have

$$\begin{aligned} \sigma(\Omega) &\propto \int_{-\infty}^{\infty} d\tau \exp[+i(\Omega - \Omega_{\text{peak}})\tau] \\ &\quad \times \exp \left[-\frac{1}{2\hbar^2} \int_0^\tau d\tau_1 \int_0^\tau d\tau_2 \langle X(\tau_1)X(\tau_2) \rangle \right] \end{aligned} \quad (\text{A205})$$

$$\approx \int_{-\infty}^{\infty} d\tau \exp \left[+i(\Omega - \Omega_{\text{peak}})\tau - \frac{\tau^2 \lambda k_B T}{\hbar^2} \right] \quad (\text{A206})$$

$$= \sqrt{\frac{\pi \hbar^2}{\lambda k_B T}} \exp \left[-\frac{(\hbar\Omega - \hbar\Omega_{\text{peak}})^2}{4\lambda k_B T} \right]. \quad (\text{A207})$$

This result should look familiar from Eq (66).

$$\sigma(\Omega) \propto \int_{-\infty}^{\infty} d\tau \exp \left[+i(\Omega - \Omega_{\text{peak}})\tau - \frac{\Delta^2 \omega^2 k_B T}{2\hbar^2} \int_0^\tau d\tau_1 \int_0^\tau d\tau_2 \cos(\omega(\tau_1 - \tau_2)) \right] \quad (\text{A210})$$

$$= \int_{-\infty}^{\infty} d\tau \exp \left[+i(\Omega - \Omega_{\text{peak}})\tau - \frac{\Delta^2 k_B T}{\hbar^2} (1 - \cos(\omega\tau)) \right]. \quad (\text{A211})$$

Now we notice that the term $\exp[-(\Delta^2 k_B T / \hbar^2) \cos(\omega\tau)]$ is periodic, and thus has a discrete Fourier expansion; the only frequencies which appear are integer multiples of the vibrational frequency ω . As a result,

$$\sigma(\Omega) = \sum_n A_n \delta(\Omega - \Omega_{\text{peak}} - n\omega). \quad (\text{A212})$$

Thus, in this limit of a single undamped mode, the absorption spectrum *does* consist of a set of sharp lines, spaced by the vibrational quanta. In order to recover the semi-classical picture, these resonances must be washed out by a combination of multiple modes (so that the discrete absorption lines become a dense forest) and some

The calculation we have done here also allows us to look more precisely at the limits to our approximation. The integral in Eq (A206) is a Gaussian integral over τ , which means that it is done exactly by the saddle point method. The characteristic time which emerges from this is

$$\delta t \sim \frac{\hbar}{\sqrt{\lambda k_B T}}. \quad (\text{A208})$$

If the typical vibrational time scales that enter into $C_X(\tau)$ are $\tau_{\text{vib}} \sim \hbar/k_B T$, then the condition for validity of our approximation becomes $\lambda \gg k_B T$. Tracing the factors through, our approximate result should be valid if the predicted width of the absorption spectrum is (in energy units) larger than $k_B T$, or roughly one percent of $\hbar\Omega_{\text{peak}}$. This is a rather gentle condition, suggesting that whenever the model of harmonic normal modes is correct, something like the saddle point approximation ought to work.

In fact, this calculation also gives us insight into another way that our semi-classical intuition from Fig 19 can fail. If, for example, there was just a single normal mode, we would have $X = gq(t)$, where $g = \omega^2 \Delta$. But if there is just this one mode, and no other degrees of freedom to suck energy out of this mode, we must have

$$\langle q(t)q(t') \rangle = \frac{k_B T}{\omega^2} \cos[\omega(t - t')], \quad (\text{A209})$$

so the integral [Eq (A205)] which defines the cross-section becomes

dissipation corresponding to a lifetime or dephasing of each individual mode.

Problem 176: Washing out resonances. Suppose that we have just a single mode, but this mode is damped so that

$$\langle q(t)q(t') \rangle = \frac{k_B T}{\omega^2} \cos[\omega(t - t')] \exp[-\gamma|t - t'|]. \quad (\text{A213})$$

If $\gamma \ll \omega$, the integral in Eq (A205) which defines the absorption cross-section is almost the integral of a period function. Thus there will be multiple saddle points, the first (the one we have considered in our semi-classical approximation) being close to $\tau = 0$, and all the others close to $\tau = 2\pi n/\omega$ for integer n . Carry out this expansion, and analyze your results. Can you see how, as $\gamma \rightarrow 0$, this sum

over saddle points gives back the discrete spectral lines? At large γ , what enforces the smooth dependence of the cross-section on Ω ? How big does γ need to be in order that we wouldn't see much hint of the vibrational resonances in the absorption spectrum? Is it possible that the vibrations are weakly damped ($\gamma \ll \omega$), but there are no visible resonances in the absorption spectrum?

Say something about the quantum treatment of the coordinate $q(t)$, and the zero-phonon lines. Maybe a word about the relation to the Moössbauer effect?

I think there is still more to say here. Notice that to make things consistent we need a quantum mechanical treatment of damping, which was a big puzzle some time back. This is also related to discussions of decoherence in more modern times. At the very least we need pointers to references. One could also note that what enters these computations are certain correlation functions of the “relevant” coordinates, and so if these correlation functions are damped (however this happens!) all will be well. Still ... an opportunity to teach some physics shouldn't be missed.

Need refs to standard text on molecular spectra; maybe old refs to solid state problem of electron-phonon couplings in impurity spectra. The idea that coupling to a bath of oscillators could describe dissipation in quantum mechanics goes back, at least, to Feynman & Vernon (1963). These ideas were revitalized by Caldeira & Leggett (1981, 1983), who were especially interested in the impact of dissipation on quantum tunneling.

Caldeira & Leggett 1981: Influence of dissipation on quantum tunnelling in macroscopic systems. AO Caldeira & AJ Leggett, *Phys Rev Lett* **46**, 211–214 (1981).

Caldeira & Leggett 1983: Quantum tunnelling in a dissipative system. AO Caldeira & AJ Leggett, *Ann Phys (NY)* **149**, 374–456 (1983).

Feynman & Vernon 1963: The theory of a general quantum system interacting with a linear dissipative system. RP Feynman & FL Vernon Jr, *Ann Phys (NY)* **24**, 118–173 (1963).

4. Cooperativity

[Be sure to talk about the specific case of hemoglobin, so we can point from Section II.A.]

To understand the statistical mechanics of cooperative interactions in the binding of multiple ligands, it is useful to start at the beginning, with the binding of a single ligand, especially since many physics students don't have much experience with problems that get categorized as “chemistry.” Suppose that we have a receptor molecule

R to which some smaller ligand molecule L can bind. For simplicity let there just be the two states, R with its binding site empty, and RL with the binding site filled by an L molecule, and let us assume that every binding event is independent, so the different receptor molecules don't interact. To study the dynamics of this system we keep track of the number of receptors in the state R and the number in state RL ; these numbers, n_R and n_{RL} , respectively, must add up to give the total number of receptors, N .

The rate at which empty sites get filled ($R \rightarrow RL$) must be proportional to the number of empty sites and to the concentration c of the ligand. The rate at which filled sites become empty should just be proportional to the number of filled sites. Thus

$$\frac{dn_{RL}}{dt} = k_+cn_R - k_-n_{RL}, \quad (\text{A214})$$

where k_+ is the rate constant for binding and k_- is the rate constant for unbinding; note that these have different units. Since $n_R + n_{RL} = N$, this becomes

$$\frac{dn_{RL}}{dt} = k_+cN - (k_- + k_+c)n_{RL}. \quad (\text{A215})$$

The equilibrium state is reached when

$$n_{RL} = N \frac{k_+c}{k_- + k_+c}. \quad (\text{A216})$$

The fraction n_{RL}/N can also be interpreted microscopically as the probability that one receptor will be the state RL ,

$$P_{RL} = \frac{k_+c}{k_- + k_+c} = \frac{c}{K + c}, \quad (\text{A217})$$

where the equilibrium constant (or “dissociation constant”) $K = k_-/k_+$.

From statistical mechanics, if we have a molecule that can be in two states, we should calculate the probability of being in these states by knowing the energy of each state and using the Boltzmann distribution. Importantly, what we mean by “state,” especially when discussing large molecules, often is a large group of microscopic configurations. Thus saying that there are two states R and RL really means that we can partition the phase space of the system into two regions, and these regions are what we label as R and RL . Then, as should be familiar, what matters is not the energy of each state but the free energy. The free energy of the state R has one component from the receptor molecule itself, F_R , plus a component from the ligand molecules in solution. In the transition $R \rightarrow RL$, the free energy of the receptor changes to F_{RL} , and the free energy of the solution changes because one molecule of the ligand is removed. The change in free energy when we add one molecule to the solution defines the chemical potential $\mu(c)$. Thus,

up to an arbitrary zero of energy, we can consider the free energy of the two states to be F_R and $F_{RL} - \mu(c)$. Then the probability of being in the state RL is given by the Boltzmann distribution,

$$P_{RL} = \frac{1}{Z} \exp\left(-\frac{F_{RL} - \mu(c)}{k_B T}\right), \quad (\text{A218})$$

where the partition function Z is given by the sum of the Boltzmann factors over both available states,

$$Z = \exp\left(-\frac{F_R}{k_B T}\right) + \exp\left(-\frac{F_{RL} - \mu(c)}{k_B T}\right). \quad (\text{A219})$$

Putting the terms together, we have

$$\begin{aligned} P_{RL} &= \frac{\exp[-(F_{RL} - \mu(c))/k_B T]}{\exp[-F_R/k_B T] + \exp[-(F_{RL} - \mu(c))/k_B T]} \\ &= \frac{e^{\mu(c)/k_B T}}{\exp[-(F_R - F_{RL})/k_B T] + e^{\mu(c)/k_B T}}. \end{aligned} \quad (\text{A220}) \quad (\text{A221})$$

Notice that the only place where the ligand concentration appears is in the chemical potential $\mu(c)$. In order for this result to be consistent with the result from analysis of the kinetics in Eq (A217), we must have $e^{\mu(c)/k_B T} \propto c$, and you may recall that when concentrations are low—as

in ideal gases, and also ideal solutions—it is a standard result that

$$\mu(c) = k_B T \ln(c/c_0), \quad (\text{A222})$$

where c_0 is some reference concentration. Then we can also identify the equilibrium constant as

$$K = c_0 \exp\left(-\frac{F_{\text{bind}}}{k_B T}\right), \quad (\text{A223})$$

where $F_{\text{bind}} = F_R - F_{RL}$ is the change in free energy when the ligand binds to the receptor.

Now suppose we have a receptor to which two ligands can bind. There are now four states, which we can think of as 00, 10, 01, and 11. If the each binding event is identical and independent, then the free energies of these states are

$$F_{00} = F_R \quad (\text{A224})$$

$$F_{01} = F_{10} = F_R - F_{\text{bind}} - \mu(c) \quad (\text{A225})$$

$$F_{11} = F_R - 2F_{\text{bind}} - 2\mu(c). \quad (\text{A226})$$

If we calculate, for example, the probability that both binding sites are occupied—i.e., that the molecule is in the state 11—we have

$$P_{11} = \frac{1}{Z} e^{-F_{11}/k_B T} \quad (\text{A227})$$

$$= \frac{\exp\left[-\frac{F_R - 2F_{\text{bind}} - 2\mu(c)}{k_B T}\right]}{\exp\left[-\frac{F_R}{k_B T}\right] + 2 \exp\left[-\frac{F_R - F_{\text{bind}} - \mu(c)}{k_B T}\right] + \exp\left[-\frac{F_R - 2F_{\text{bind}} - 2\mu(c)}{k_B T}\right]} \quad (\text{A228})$$

$$= \frac{(c/K)^2}{1 + 2(c/K) + (c/K)^2} = \left(\frac{c}{c + K}\right)^2. \quad (\text{A229})$$

Thus, the probability of both sites being occupied is just the square of the probability that a single binding site will be occupied, as in Eq (A217). This makes sense, because we assumed that binding to the two sites were independent events.

Problem 177: Counting bound molecules. Rather than counting the fraction of molecules in the doubly bound state, count the number of ligands bound. Show that this is just $2 \times c/(c + K)$, and explain why.

In fact, in many cases we see that binding of multiple ligands to a protein molecule are not independent events. As a start, let's suppose that we again have two binding sites, but the doubly bound state is stabilized (for as yet unspecified reasons) by an extra energy Δ . Then if we calculate the fraction of binding sites occupied, we have

$$f = \frac{1}{2} [P_{01} + P_{10} + 2P_{11}] \quad (\text{A230})$$

$$= \frac{1}{2} \frac{2 \exp \left[-\frac{F_R - F_{\text{bind}} - \mu(c)}{k_B T} \right] + 2 \exp \left[-\frac{F_R - 2F_{\text{bind}} - 2\mu(c) - \Delta}{k_B T} \right]}{\exp \left[-\frac{F_R}{k_B T} \right] + 2 \exp \left[-\frac{F_R - F_{\text{bind}} - \mu(c)}{k_B T} \right] + \exp \left[-\frac{F_R - 2F_{\text{bind}} - 2\mu(c) - \Delta}{k_B T} \right]} \quad (\text{A231})$$

$$= \frac{c/K + J(c/K)^2}{1 + 2(c/K) + J(c/K)^2}, \quad (\text{A232})$$

where $J = \exp(\Delta/k_B T)$. Results are shown in Fig 168. We see that, as the interaction energy increases, the binding sites can be occupied at lower concentration, but more importantly the steepness of the “switch” from empty to full sites is more abrupt. This abruptness is the signature of cooperativity.




FIG. 168 Cooperative binding, with two binding sites that interact. Lines show the predicted fraction of binding sites vs. concentration, for different values of the interaction energy Δ ; from Eq (A232).

The classic example is the oxygen binding protein hemoglobin in our blood. We now know that hemoglobin has four protein subunits, each of which has an iron atom which can bind one oxygen molecule. **[Might be good to show some figures from Hb!]** As Hill recognized in the early part of the twentieth century, the fraction of sites with bound oxygen behaves more nearly as if all four molecules had to bind together, so that

$$f = \frac{c^n}{c^n + K^n}, \quad (\text{A233})$$

with $n = 4$; this is still called a “Hill function” in many contexts. As shown in Fig 169, the binding is now sigmoidal, or more nearly switch like at larger n . Because the natural quantity in statistical mechanics is the chemical potential and not the concentration, things look simpler on a logarithmic concentration axis. Cooperative




FIG. 169 Cooperative binding, in the Hill model. Blue lines show the predicted fraction of binding sites vs. concentration when binding to each site in independent. Green lines show the case of cooperative binding to four sites, as described by the Hill model in Eq (A233), with $n = 4$. At left, a linear concentration scale; at right, a logarithmic scale.

binding corresponds to a steeper slope on these logarithmic plots. This is clear for the Hill function, where we can see that

$$\frac{dF}{dc} = \frac{n}{c} F(1 - F), \quad (\text{A234})$$

and hence

$$\left. \frac{dF}{d \ln c} \right|_{F=1/2} = \frac{n}{4}, \quad (\text{A235})$$

so the slope is a direct measure of the number of molecules forced to bind simultaneously. Few real systems are described exactly by the Hill model, but it’s a good approximation. We should also appreciate the power of Hill’s intuition, in seeing the connection of the sigmoidal binding curves to the number of protein subunits even before much was known about these molecules.

The Hill model suggests that there is some direct interaction between binding events that causes all of the ligands to bind (or not to bind) simultaneously, which we can think of as a limiting case of the model above, with $\Delta \rightarrow \infty$. In some cases, including hemoglobin,

there is little evidence for such a direct interaction. An alternative is to imagine that the whole system can be in two states. In the case of hemoglobin these came to be called ‘relaxed’ (R) and ‘tense’ (T), but in other systems there natural choices; for example, in the case of the ion channels in rod cells that open in response to binding of cGMP, the two states might simply be the open and closed states of the channel, as in Fig 170. To continue with this example, the channel can bind one, two or three molecules of cGMP. If all the binding sites are empty, the free energies of the two states are F_{open} and F_{closed} . Given that the channel is closed, the binding of a single cGMP molecule lowers the energy by an amount $F_{\text{closed}}^{\text{bind}}$, but in addition this takes one molecule out of the solution and hence the free energy of the system also goes down by μ , the chemical potential. So the total free energy of the state with the channel closed and one molecule bound is

$$F_{\text{closed}}(1) = F_{\text{closed}} - F_{\text{closed}}^{\text{bind}} - \mu \quad (\text{A236})$$

$$= F_{\text{closed}} - F_{\text{closed}}^{\text{bind}} - k_B T \ln(c/c_0) \quad (\text{A237})$$

$$= F_{\text{closed}} - k_B T \ln\left(\frac{c}{K_{\text{closed}}}\right), \quad (\text{A238})$$

and similarly for the open state,

$$F_{\text{open}}(1) = F_{\text{open}} - k_B T \ln\left(\frac{c}{K_{\text{open}}}\right). \quad (\text{A239})$$

The important point is that the binding energies to the open and closed states are different. By detailed balance, this means that, as the cGMP molecules bind, they will shift the equilibrium between open and closed. The two state model was proposed by Monod, Wyman and Changeaux. They made the simplifying assumption that the only source of cooperativity among the binding events was this shifting of equilibria, so that if the target protein is in one state, each binding event remains independent, and then the free energies work out as in Fig 170.

Problem 178: Cooperativity in the MWC model. Show that the model in Fig 170 is equivalent to the statement that the free energy difference between open and closed states has a term proportional to the number of cGMP molecules bound. What is this proportionality constant in terms of the other parameters? Can you explain the connection between these two points of view on the model?

It’s a useful exercise to work out the statistical mechanics of the MWC model. The partition function has two classes of terms, coming from the two states of the protein. In each state, we have to sum over the occupied and unoccupied states of each binding site, but this is




FIG. 170 A model for binding of cGMP to the channels in rod cells. Cooperativity arises not from direct interactions among the cGMP molecules but rather because binding of each molecule contributes to stabilizing a different structure of the channel protein. In this case the two structures are just the open and closed states.

relatively easy because the sites are independent. In the notation of Fig 170, we find

$$Z = Z_{\text{open}} + Z_{\text{closed}} \quad (\text{A240})$$

$$Z_{\text{open}} = \exp\left(-\frac{F_{\text{open}}}{k_B T}\right) \left(1 + \frac{c}{K_{\text{open}}}\right)^n \quad (\text{A241})$$

$$Z_{\text{closed}} = \exp\left(-\frac{F_{\text{closed}}}{k_B T}\right) \left(1 + \frac{c}{K_{\text{closed}}}\right)^n, \quad (\text{A242})$$

where in the case of the cGMP-gated channels, $n = 3$. The probability of being in the open state is then

$$P_{\text{open}} = \frac{Z_{\text{open}}}{Z_{\text{open}} + Z_{\text{closed}}} \quad (\text{A243})$$

$$= \frac{(1 + c/K_{\text{open}})^n}{(1 + c/K_{\text{open}})^n + L(1 + c/K_{\text{closed}})^n} \quad (\text{A244})$$

where $k_B T \ln L = F_{\text{open}} - F_{\text{closed}}$ is the free energy difference between open and closed states in absence of ligand binding. In the limit that binding is much stronger to the open state, $K_{\text{open}} \ll K_{\text{closed}}$, this simplifies,

$$P_{\text{open}} = \frac{(1 + c/K_{\text{open}})^n}{L + (1 + c/K_{\text{open}})^n} \quad (\text{A245})$$

$$= \frac{1}{1 + \exp[\theta - n \ln(1 + c/K_{\text{open}})]}, \quad (\text{A246})$$

where $\theta = \ln L$. This is similar to the Hill model, but a little different in detail. Distinguishing the models from the equilibrium data alone is difficult, but clearly the MWC model predicts that binding has an extra kinetic step in which the protein makes the transition between its two states; if we are lucky we can ‘catch’ the system

after the first ligand molecules have bound but before this change in protein structure. Indeed, such experiments were critical in understanding the mechanism of cooperativity in hemoglobin.

Problem 179: Details of the MWC model. Fill in the steps to Eqs (A240–A242). Then, compare the Hill model with MWC. Show that for $c \gg K_{\text{open}}$, Eq (A246) reduces to Eq (A233). What about at $c \ll K_{\text{open}}$? The MWC model, even in the limit $K_{\text{open}} \ll K_{\text{closed}}$, has one more parameter than the Hill model; what does this freedom mean for the class of functions that the MWC model can realize?

In many systems, it is not just a single class of ligands that binds. For hemoglobin itself, changes in pH, which presumably result in binding and unbinding of protons, change the way in which oxygen binds.⁹⁵ For enzymes—proteins that catalyze a chemical reaction—it is not just the substrate which binds and is chemically altered, but other molecules bind as well and alter the activity of the enzyme. It is important that these ‘other molecules’ are binding at other sites, not directly interfering with substrate binding in enzymes or oxygen binding in hemoglobin. From the Greek for “other site,” these effects are called “allosteric,” and the MWC model gives a framework for a much more general view of allostery. In this view, *all* binding events are independent, but with binding energies that depend on the overall state of the target protein. In this way, all binding events can shift the R/T equilibrium.

Maybe another problem? There should be a figure with data! Tell the story about Perutz? Need to flesh out the text to match references. Put something about protein/DNA interactions here?

The classic paper on “Hill functions” for cooperative binding is Hill (1910). There is some suggestion that Hill might have been the first to derive the simpler description of independent binding, often called the “Langmuir” isotherm; for this and more related history as seen through the lens of drug–receptor interactions, see Colquhoun (2006). The MWC model is due to Monod et al (1965), and a contemporary, competing model is due to Koshland et al (1966). Late in his life, Perutz (1990) provided some perspective on his long adventure with hemoglobin. A key step in understanding was to show, convincingly, that there really is no direct interaction between the binding sites, and the cooperativity was mediated entirely by the shifting equilibrium between the R and T states (Shulman et al 1975). The MWC model leaves open the question of where the energy for cooperativity is stored in the molecule; for a hypothesis very much ahead of its time, see Hopfield (1973).

Colquhoun 2006: The quantitative analysis of drug–receptor interactions: A short history. D Colquhoun, *Trends Pharm Sci* **27**, 149–157 (2006).

Hill 1910: The possible effects of the aggregation of the molecules of haemoglobin on its dissociation curves. AV Hill, *J Physiol (Lond)* **40**, Suppl iv–vii (1910).

Hopfield 1973: Relation between structure, co-operativity and spectra in a model of hemoglobin action. JJ Hopfield, *J Mol Biol* **77**, 207–222 (1973).

Koshland et al 1966: Comparison of experimental binding data and theoretical models in proteins containing subunits. DE Koshland Jr, G Némethy & D Filmer, *Biochemistry* **5**, 365–385 (1966).

Monod et al 1965: On the nature of allosteric transitions: A plausible model. J Monod, J Wyman & JP Changeux, *J Mol Biol* **12**, 88–118 (1965).

Perutz 1990: *Mechanisms of Cooperativity and Allosteric Regulation in Proteins*. MF Perutz (Cambridge University Press, Cambridge, 1990).

Shulman et al 1975: Allosteric interpretation of haemoglobin properties. RG Shulman, JJ Hopfield & S Ogawa, *Q Rev Biophys* **8**, 325–420 (1975).

The idea that interesting things happen in “cooperative” events in larger and larger collections of interacting subunits probably occurred to many people who thought in terms of statistical mechanics (Thompson 1972), and it seems to get rediscovered periodically (Bray et al 1998, Duke & Bray 1999, Duke et al 2001). Statistical mechanics models for interactions among binding events also play a key role in thinking about protein/DNA interactions and the regulation of gene expression (Bintu et al 2005a,b, Kinney et al 2010).

Bintu et al 2005a: Transcriptional regulation by the numbers: models. L Bintu, NE Buchler, HG Garcia, U Gerland, T Hwa, J Kondev & R Phillips, *Curr Opin Gene Dev* **15**, 116–124 (2005).

Bintu et al 2005b: Transcriptional regulation by the numbers: applications. L Bintu, NE Buchler, HG Garcia, U Gerland, T Hwa, J Kondev, T Kuhlman & R Phillips, *Curr Opin Gene Dev* **15**, 125–136 (2005).

Bray et al 1998: Receptor clustering as a cellular mechanism to control sensitivity. D Bray MD Levin & CJ Morton–Firth, *Nature* **393**, 85–88 (1998).

Duke & Bray 1999: Heightened sensitivity of a lattice of membrane receptors. TAJ Duke & D Bray, *Proc Nat'l Acad Sci (USA)* **96**, 10104–10108 (1999).

Duke et al 2001: Conformational spread in a ring of proteins: A stochastic view of allostery. TAJ Duke, N Le Novère & D Bray, *J Mol Biol* **308**, 541–553 (2001).

Kinney et al 2010: Using deep sequencing to characterize the biophysical mechanism of a transcriptional regulatory sequence. JB Kinney, A Murugan, CG Callan Jr & EC Cox, *Proc Nat'l Acad Sci (USA)* **107**, 9158–9163 (2010).

Thompson 1971: *Mathematical Statistical Mechanics* CJ Thompson (Macmillan, New York, 1972).

⁹⁵ This is the “Bohr effect,” after Christian Bohr, Niels’ father.

5. X-ray diffraction and biomolecular structure

The first detailed experimental information about the structure of biological molecules came from X-ray diffraction measurements. We recall that if a particle scatters from a sample, shifting its energy by $\hbar\omega$ and its momentum by $\hbar\vec{q}$, then the amplitude for this scattering event must be proportional to the (\vec{q}, ω) spatiotemporal Fourier component of the relevant density in the sample. For an electromagnetic wave what matters is (roughly) the charge density. Thus, the cross-section for elastic ($\omega = 0$) scattering is

$$\sigma(\vec{q}) \propto \left| \int d^3x e^{i\vec{q}\cdot\vec{x}} \rho(\vec{x}) \right|^2. \quad (\text{A247})$$

It is useful to have in mind the geometry [ref to a Fig]. If the X-ray photons approach the sample collimated along the \hat{x} axis, they have an initial wavevector $\vec{k}_0 = k\hat{x}$, where as usual $k = 2\pi/\lambda$, with λ the wavelength. If they emerge with a final wavevector \vec{k}_f at an angle θ relative to the \hat{x} axis, then $\vec{q} \equiv \vec{k}_f - \vec{k}_0$, and the magnitude of the scattering vector (or, up to a factor \hbar , momentum transfer) is

$$|\vec{q}| = |\vec{k}_f - \vec{k}_0| \quad (\text{A248})$$

$$= \sqrt{|\vec{k}_f - \vec{k}_0|^2} \quad (\text{A249})$$

$$= \sqrt{|\vec{k}_f|^2 - 2\vec{k}_f \cdot \vec{k}_0 + |\vec{k}_0|^2} \quad (\text{A250})$$

$$= \sqrt{k^2 - 2k^2 \cos \theta + k^2} \quad (\text{A251})$$

$$= \sqrt{2k^2(1 - \cos \theta)} = 2k \sin(\theta/2). \quad (\text{A252})$$

Thus scattering by a small angle corresponds to a small momentum transfer. The classic results about X-ray diffraction concern the case where the density profile is periodic, as in a crystal. If the periodicity corresponds to displacement by d (let's think along one dimension, for the moment), then the density can be expressed as a discrete Fourier series, which means [from Eq (A247)] that $\sigma(\vec{q})$ will have delta functions at $|\vec{q}| = 2\pi n/d$, with n an integer. Combining this with Eq (A252), we find the angles which satisfy the ‘‘Bragg condition,’’

$$2\pi n/d = (4\pi/\lambda) \sin(\theta/2) \Rightarrow \sin(\theta/2) = n\lambda/2d. \quad (\text{A253})$$

[I think this is a bit off the usual way of stating the condition (2's in the wrong places); check!]

The first great triumph of X-ray diffraction in elucidating the structure of biological molecules came with the structure of DNA. This is an often told, and often distorted, piece of scientific history. Watson and Crick predicted the structure of DNA by arguing that a few key facts about the molecule, when combined with the rules of chemical bonding, were enough to suggest an interesting structure that would have consequences for the mechanisms of genetic inheritance. It was known that

DNA was composed of four different kinds of nucleotide bases: adenine (A), thymine (A), guanine (G) and cytosine (C). Importantly, Chargaff had surveyed the DNA of many organisms and shown that while the ratios of A to G, for example, vary enormously, the ratios A/T and C/G do not. Watson and Crick realized that the molecular structures of the bases are such that A and T can form favorable hydrogen bonds, as can C and G; further, the resulting hydrogen bonded base pairs are the same size, and thus could fit comfortably into a long polymer, as shown in Fig 171. Piling on top of one another, the base pairs would also experience a favorable ‘‘stacking’’ interaction among the π -bonded electrons in their rings. Finally, if one looks carefully at all the bond angles where the planar bases connect to the sugars and phosphate backbone, each successive base pair must rotate relative to its neighbor, and although there is some flexibility the favored angle was predicted to be $2\pi/10$ radians, or 36° .




FIG. 171 The structure of DNA, from Watson and Crick (1953b). At left, the polymeric pattern of bases, sugars and phosphates, and the famous double helix. At right, the pairings A/T and G/C, illustrating the similar sizes of the correct pairs. Note that the donor/acceptor pattern of hydrogen bonds discriminates against the incorrect A/C and G/T pairings.

Quite independently of his collaboration with Watson, Crick has been interested in the structure of helical molecules, and in the X-ray diffraction patterns that they should produce. Thus, when Watson and Crick realized that the structure of DNA might be a helix, they were in a position to calculate what the diffraction patterns should look like, and thus compare with the data emerging from the work of Franklin, Wilkins and collaborators. So, let's look at the theory of diffraction from a helix.

It's best to describe a helix in cylindrical coordinates: z along the axis of the helix, r outward from its center, and an angle ϕ around the axis. Helical symmetry is the statement that translations along z are equivalent

to rotations of the angle ϕ . Thus, a continuous helical structure would have the property that

$$\rho(z, r, \phi) = \rho(z + d, r, \phi + 2\pi d/\ell), \quad (\text{A254})$$

for any displacement d , where ℓ is the displacement corresponding to a complete rotation. For a discrete helical structure, the same equation is true, but only for values of d that are integer multiples of a fundamental spacing d_0 .

For the continuous helix, the dependence on the two variables z and ϕ really collapses to a dependence on one combined variable,

$$\rho(z, r, \phi) = g(r, \phi - 2\pi z/\ell). \quad (\text{A255})$$

We know that any function of angle can be expanded as a discrete Fourier series,

$$f(\phi) = \sum_{n=-\infty}^{\infty} \tilde{f}_n e^{-in\phi}, \quad (\text{A256})$$

so in this case we have

$$\rho(z, r, \phi) = \sum_{n=-\infty}^{\infty} \tilde{g}_n(r) e^{-in(\phi - 2\pi z/\ell)}. \quad (\text{A257})$$

$$\int d^3x e^{i\vec{q}\cdot\vec{x}} \rho(\vec{x}) = \int_{-\infty}^{\infty} dz \int_0^{\infty} dr r \int_0^{2\pi} d\phi e^{iq_z z} \sum_{n=-\infty}^{\infty} J_n(q_{\perp} r) e^{in\phi} \sum_{m=-\infty}^{\infty} \tilde{g}_m(r) e^{-im(\phi - 2\pi z/\ell)} \quad (\text{A262})$$

$$= \sum_{n,m=-\infty}^{\infty} \int_{-\infty}^{\infty} dz e^{iq_z z} e^{-i2\pi m z/\ell} \int_0^{\infty} dr r J_n(q_{\perp} r) \tilde{g}_m(r) \int_0^{2\pi} e^{in\phi} e^{-im\phi}. \quad (\text{A263})$$

We see that the integral over ϕ forces $m = n$, and the integral over z generates delta functions at $q_z = 2\pi n/\ell$. Thus, for a continuous helix we expect that the X-ray scattering cross section will behave as

$$\sigma(q_z, q_{\perp}) \propto \sum_{n=-\infty}^{\infty} \delta(q_z - 2\pi n/\ell) \left| \int_0^{\infty} dr r J_n(q_{\perp} r) \tilde{g}_n(r) \right|^2. \quad (\text{A264})$$

In particular, if most of the density sits at a distance R from the center of the helix (which is not a bad approximation for DNA, since the phosphate groups have much more electron density than the rest of the molecule), then

$$\sigma(q_z, q_{\perp}) \sim \sum_{n=-\infty}^{\infty} \delta(q_z - 2\pi n/\ell) \left| J_n(q_{\perp} R) \right|^2. \quad (\text{A265})$$

Equation (A265) is telling us that diffraction from a helix generates a series of “layer lines” at $q_z = 2\pi n/\ell$, and from their spacing we should be able to read off the “pitch” of the helix, the distance ℓ along the \hat{z} axis corresponding to a complete turn. Further, if we look along

Our task is to compute

$$\int d^3x e^{i\vec{q}\cdot\vec{x}} \rho(\vec{x}). \quad (\text{A258})$$

In cylindrical coordinates, we can write $\vec{q} = (q_z \hat{z}, \vec{q}_{\perp})$, so that $\vec{q}\cdot\vec{x} = q_z z + q_{\perp} r \cos \phi$, where we choose the origin of the angle ϕ to make things simple and $q_{\perp} = |\vec{q}_{\perp}|$. Thus we have

$$e^{i\vec{q}\cdot\vec{x}} = e^{iq_z z} e^{iq_{\perp} r \cos \phi} \quad (\text{A259})$$

$$= e^{iq_z z} \sum_{n=-\infty}^{\infty} J_n(q_{\perp} r) e^{in\phi}, \quad (\text{A260})$$

where [check the conventions for the definition of the Bessel function!]

$$J_n(u) = \int_0^{2\pi} \frac{d\phi}{2\pi} e^{-in\phi} e^{iu \cos \phi} \quad (\text{A261})$$

are Bessel functions. Putting Eq (A260) together with the consequences of helical symmetry in Eq (A257), we have

a single layer line, we should see an intensity varying as




FIG. 172 Diffraction from continuous (left) and discrete (right) helices; Holmes (1998).




FIG. 173 The justly famous photograph 51, showing the diffraction from DNA molecules pulled into a fiber, from Franklin & Gosling (1953).

$\sim |J_n(q_\perp R)|^2$. What is important here about the Bessel functions is that for small q_\perp we have $J_n(q_\perp R) \propto (q_\perp R)^n$, and the first peak of the n^{th} Bessel function occurs at a

point roughly proportional to n . The resulting pattern is shown schematically in Fig 172.

Problem 180: Bessel functions. Verify the statements about Bessel functions made above, in enough detail to understand the diffraction patterns shown in Fig 172.

Let's see what happens when we move from the continuous to the discrete helix. To keep things simple, suppose that all the density indeed is concentrated at a distance R from the center of the helix, so that

$$\rho(\vec{x}) = \frac{1}{R} \delta(r - R) \sum_n \delta(z - nd_0) \delta(\phi - n\phi_0), \quad (\text{A266})$$

where the rotation from one element to the next $\phi_0 = 2\pi d_0/\ell$; notice that we don't really require ℓ/d_0 to be an integer. Now we have

$$\int d^3x e^{i\vec{q}\cdot\vec{x}} \rho(\vec{x}) = \int_{-\infty}^{\infty} dz \int_0^{\infty} dr r \int_0^{2\pi} d\phi e^{iq_z z} \sum_{n=-\infty}^{\infty} J_n(q_\perp r) e^{in\phi} \frac{1}{R} \delta(r - R) \sum_{m=-\infty}^{\infty} \delta(z - md_0) \delta(\phi - m\phi_0) \quad (\text{A267})$$

$$= \sum_{n=-\infty}^{\infty} J_n(q_\perp R) \sum_{m=-\infty}^{\infty} \int_{-\infty}^{\infty} dz \delta(z - md_0) e^{iq_z z} \times \int_0^{2\pi} d\phi \delta(\phi - m\phi_0) e^{in\phi} \quad (\text{A268})$$

$$= \sum_{n=-\infty}^{\infty} J_n(q_\perp R) \sum_{m=-\infty}^{\infty} e^{im(n\phi_0 + q_z d_0)} \quad (\text{A269})$$

$$= \sum_{n=-\infty}^{\infty} J_n(q_\perp R) \sum_{m=-\infty}^{\infty} \delta(n\phi_0 + q_z d_0 - 2\pi m) \quad (\text{A270})$$

$$\propto \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} J_n(q_\perp R) \delta(q_z + 2\pi n/\ell - 2\pi m/d_0). \quad (\text{A271})$$

Thus the discrete helix involves a double sum of terms. If we set $m = 0$ we have the results for the continuous helix. But the sum over $m \neq 0$ causes the whole “X” pattern of the continuous helix to be repeated with centers at $(q_z = 2\pi m/d_0, q_\perp = 0)$; the line $q_\perp = 0$ is often called the meridian, and so the extra peaks centered on $(q_z = 2\pi m/d_0, q_\perp = 0)$ are called meridional reflections. All of this is shown in Fig 172. Just as the spacing of the layer lines allows us to measure the helical pitch ℓ , the spacing of the meridional reflections allows us to measure the spacing d_0 between discrete elements along the helix.

At this point you know what Watson and Crick knew [maybe put in the precise dates of these events, from Watson's memoir]. They had a theory of what the structure

should be, and almost certainly they had already realized the implications of this structure, as they remarked in their first paper “It has not escaped our notice that the specific pairing we have postulated immediately suggests a possible copying mechanism for the genetic material.” They also knew that if the structure was as they had theorized, then the diffraction pattern should display a number of key signatures—the regularly spaced layer lines, the “X” arrangement of their intensities, and the meridional reflections—that would provide both qualitative and quantitative confirmation of the theory. Thus you should be able to imagine their excitement when they saw the clean X-ray diffraction pattern from hydrated DNA, the famous photograph 51 taken by Ros-

alind Franklin, Fig 173. As far as one could tell, the proposed structure was right.

Problem 181: Discrete helices, more generally. Show that most of what was said above can be generalized to an arbitrary discrete helix, without assuming that the density is concentrated at $r = R$. That is, use only the symmetry defined by Eq (A254) for $d = nd_0$.

Problem 182: Fibers vs. crystals. We have discussed the diffraction from a helix as if there were just one molecule, and we have not been very precise about the difference between amplitudes and intensities. Show that if there are many helices, all with their \hat{z} axes aligned but with random positions and orientations in the $\hat{x} - \hat{y}$ plane, then the diffraction intensity from the ensemble of molecules depends only on the structure of the individual helices, and that all directions for the vector \vec{q}_\perp are equivalent.

It is crucial to appreciate that, contrary to what is often said in textbooks, it was not possible to “determine” the structure of DNA by looking at diffraction patterns like those in Fig 173. On the other hand, if you thought you knew the structure, you could predict the diffraction pattern—in the regime where it could be measured—and see if you got things right. This difference between experiments that support a theory, or which find something that a theory tells us must exist, and experiments that “discover” something unexpected or genuinely unknown is an incredibly important distinction, often elided.

So much has been written about this moment in scientific history that it would be irresponsible not to pause and reflect. On the other hand, I am not a historian. So let me make just a few observations. Most importantly, I think, the story of the DNA structure combines so many themes in our understanding of science and society (separately and together) that it has an almost mythical quality, and as with the ancient myths everyone can see something that connects to their own concerns. There is the enormous issue of gender in the scientific community, something for which we hardly even had a vocabulary until decades after the event. There are the personalities of all the individuals, both as they were in 1953 and as they developed in response to the world-changing discovery in which they participated. There is the tragedy of Franklin’s early death. There is the competition between Cambridge and London, and the impact of an American interloper on these very British social structures. Finally, there are issues that are more purely about the science, such as the interaction between theory and experiment, physics and biology. **We could wander in this part of history for a long time. I need to come back and see what is essential, and what can be skipped. For now, let’s move on.**

In order to actually **determine** the structure of a large molecule by X-ray diffraction, we need to form crystals

of those molecules. Crystals of a protein are not like crystals of salt or even small molecules. They are quite soft, and contain quite a lot of water. The bonds between proteins, for example, in a crystal are much weaker than the bonds that hold each protein together. On the one hand this makes growing and handling the crystals quite difficult. On the other hand, it means that the internal structures of the protein in the crystal is more likely to be typical of its structure when free in solution.

We recall that being a crystal in three dimensions means that there are vectors \vec{a} , \vec{b} , and \vec{c} such that the density is the same if we translate by integer combinations of these vectors,

$$\rho(\mathbf{x}) = \rho(\mathbf{x} + n\vec{a} + m\vec{b} + k\vec{c}). \quad (\text{A272})$$

This means that the density can be expanded into a Fourier series,

$$\rho(\mathbf{x}) = \sum_{knm} \tilde{\rho}_{knm} \exp \left[i(k\vec{G}_a + n\vec{G}_b + m\vec{G}_c) \cdot \vec{x} \right], \quad (\text{A273})$$

where the \vec{G}_i are the “reciprocal lattice vectors.” As a result, the X-ray scattering cross-section is a set of delta functions or “Bragg peaks,”

$$\sigma(\vec{q}) \propto \sum_{knm} |\tilde{\rho}_{knm}|^2 \delta(\vec{q} - k\vec{G}_a - n\vec{G}_b - m\vec{G}_c). \quad (\text{A274})$$

Problem 183: Details of diffraction. Fill in the details leading to Eq (??), including the relationship between the reciprocal lattice vectors \vec{G}_i and the real lattice vectors \vec{a} , \vec{b} , and \vec{c} .

Even if we can make a perfect measurement of $\sigma(\vec{q})$, we only learn about the magnitudes of the Fourier coefficients, $|\tilde{\rho}_{knm}|^2$, and this isn’t sufficient to reconstruct the density $\rho(\vec{x})$. This is called the phase problem. For small structures it is not such a serious problem, since the constraint that $\rho(\vec{x})$ has to be built out of discrete atoms allows us to determine the positions of the atoms from the diffraction pattern. But for a protein, with thousands of atoms in each unit cell of the crystal, this is hopeless.

The phase problem was solved experimentally through the idea of “isomorphous replacement.” Suppose that we could attach to the each molecule in the crystal one or more very heavy atoms, in well defined (but unknown) positions. If we can do this without disrupting the packing of the molecules into the crystal, then the positions of the Bragg peaks will not change, but their intensities will. If we can approximate the density profiles of the

heavy atoms as delta functions (which should be right unless we look at very large $|\vec{q}|$), then

$$|\tilde{\rho}_{knm}|^2 \rightarrow \left| \rho_{knm} + \sum_{\mu} Z_{\mu} e^{i\vec{q}_{knm} \cdot \vec{x}_{\mu}} \right|^2, \quad (\text{A275})$$

where $\vec{q}_{knm} = k\vec{G}_a - n\vec{G}_b - m\vec{G}_c$, Z_{μ} is the charge of the μ^{th} heavy atom and \vec{x}_{μ} is its position. In the simple case of one added heavy atom, we can choose coordinates so that its position is at the origin, and then it should be clear that the change in intensity on adding the heavy atom is directly sensitive to the value of $\cos \phi_{knm}$, where ϕ_{knm} is the phase of the complex number ρ_{knm} . Thus, one needs at least two different examples of adding heavy atoms to determine the phases unambiguously.

Do we need to say more here? Show in detail how two replacements determines the phase? Give a problem? I honestly don't know if one has to rely on absolute measurements, as one might think naively from the equations ... check!! Say something about other approaches to the phase problem.

The density really consists of discrete blobs corresponding to atoms, and—if we can look at sufficiently high resolution—additional density in the bonds between atoms. For the moment let's think just about the atoms. Then the density has the form

$$\rho(\vec{x}) \approx \sum_{\mu} f_{\mu} \delta(\vec{x} - \vec{x}_{\mu}), \quad (\text{A276})$$

where \vec{x}_{μ} is the position of the μ^{th} atom and f_{μ} is an effective charge or scattering density associated with that atom. Thus the scattering cross-section behaves as

$$\sigma(\vec{q}) \sim \sum_{\mu\nu} f_{\mu} f_{\nu} e^{i\vec{q} \cdot (\vec{x}_{\mu} - \vec{x}_{\nu})}. \quad (\text{A277})$$

Importantly, the positions of atoms fluctuate. The time scale of these fluctuations typically is much shorter than the time scale of the experiment, so we will see an average,

$$\sigma(\vec{q}) \sim \left\langle \sum_{\mu\nu} f_{\mu} f_{\nu} e^{i\vec{q} \cdot (\vec{x}_{\mu} - \vec{x}_{\nu})} \right\rangle. \quad (\text{A278})$$

If we assume that the fluctuations in position are Gaussian around some mean, then

$$\begin{aligned} \sigma(\vec{q}) &\sim \left\langle \sum_{\mu\nu} f_{\mu} f_{\nu} e^{i\vec{q} \cdot (\vec{x}_{\mu} - \vec{x}_{\nu})} \right\rangle \\ &\equiv \sum_{\mu\nu} f_{\mu} f_{\nu} \left\langle e^{i\vec{q} \cdot \vec{r}_{\mu\nu}} \right\rangle \end{aligned} \quad (\text{A279})$$

$$\sim \sum_{\mu\nu} f_{\mu} f_{\nu} e^{i\vec{q} \cdot \vec{r}_{\mu\nu}} e^{-\frac{1}{2}|\vec{q}|^2 \langle (\delta\vec{r}_{\mu\nu})^2 \rangle}, \quad (\text{A280})$$

where $\vec{r}_{\mu\nu} = \vec{x}_{\mu} - \vec{x}_{\nu}$, and for simplicity we assume that the fluctuations are isotropic. What we see is that

the scattering intensity at \vec{q} is attenuated relative to what we expect from a fixed structure, by an amount $e^{-\frac{1}{2}|\vec{q}|^2 \langle (\delta\vec{r}_{\mu\nu})^2 \rangle}$. These are called the Debye–Waller factors. Thus, although X-ray diffraction is a static method, it is sensitive to dynamical fluctuations in structure, although it can't really distinguish between dynamics and static disorder in the crystal.

Need to come back and see what else needs to be said, given what we need in the main text. Is it worth talking about other methods, such as EM and NMR? The motifs of protein structure? ... not sure what we need or want.

You should read the classic trio of papers on DNA structure, which appeared one after the other in the April 25, 1953 issues of *Nature*: Watson & Crick (1953a), Wilkins et al (1953) and Franklin & Gosling (1953). The foundations of helical diffraction theory had been given just a year before by Cochran et al (1952); a brief account is given by Holmes (1998). The astonishing realization that the structure of DNA implies a mechanism for the transmission of information from generation to generation was presented by Watson & Crick (1953b). It is especially interesting to read their account of the questions raised by their proposal, and to see how their brief list became the agenda for the emerging field of molecular biology over the next two decades. The rest is history, as the saying goes, so you should read at least one history book (Judson 1979).

Cochran et al 1952: The structure of synthetic polypeptides. I. The transform of atoms on a helix. W Cochran, FHC Crick & V Vand, *Acta Cryst* **5**, 581–586 (1952).

Franklin & Gosling 1953: Molecular configuration in sodium thymonucleate. RE Franklin & RG Gosling. *Nature* **171**, 740–741 (1953).

Holmes 1998: Fiber diffraction. KC Holmes, <http://www.mpimf-heidelberg.mpg.de/~holmes/fibre/branden.html> (1998).

Judson 1979: *The Eighth Day of Creation* HF Judson (Simon and Schuster, New York, 1979).

Watson & Crick 1953a: A structure for deoxyribose nucleic acid. JD Watson & FHC Crick, *Nature* **171**, 737–739 (1953).

Watson & Crick 1953b: Genetical implications of the structure of deoxyribonucleic acid. JD Watson & FHC Crick, *Nature* **171**, 964–967 (1953).

Wilkins et al 1953: Molecular structure of deoxypentose nucleic acids. MHF Wilkins, AR Stokes & HR Wilson, *Nature* **171**, 738–740 (1953).

Need classic refs about protein structure and crystallography; more if we do more.

6. Berg and Purcell, revisited

In the spirit of Berg and Purcell's original discussion, the simplest example of noise in a chemical system is just to consider the fluctuations in concentration as seen in a small volume. To treat this rigorously, let's remember

that diffusion in and out of the volume keeps the system at equilibrium. Thus, fluctuations in the concentration should be just like Brownian motion or Johnson noise. What's a little odd is that while the strength of Johnson noise is proportional to the absolute temperature, our intuition about counting molecules and the \sqrt{N} rule doesn't seem to have a place for T . So, let's see how this works.⁹⁶

If we measure the current flowing across a resistor in thermal equilibrium at temperature T , we will find a noise in the current that has a spectral density $S_I = 2k_B T/R$, where R is the resistance. More generally, if we measure between two points in a circuit, and find a frequency dependent, complex impedance $\tilde{Z}(\omega)$, then the spectral density of current noise will be

$$S_I(\omega) = 2k_B T \operatorname{Re} \left[\frac{1}{\tilde{Z}(\omega)} \right], \quad (\text{A281})$$

where Re denotes the real part. In a mechanical system it is more natural to talk about positions and forces instead of currents and voltages. Now if we measure the position and apply a force, we have a "mechanical response function" $\tilde{\alpha}(\omega)$ analogous to the (inverse) impedance,

$$\tilde{x}(\omega) = \tilde{\alpha}(\omega) \tilde{F}(\omega), \quad (\text{A282})$$

where $\tilde{x}(\omega)$ is the Fourier component⁹⁷ of $x(t)$,

$$x(t) = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} e^{-i\omega t} \tilde{x}(\omega), \quad (\text{A283})$$

and similarly for the force $\tilde{F}(\omega)$. The analog of Eq (A281) for Johnson noise is that the fluctuations in position x have a spectral density

$$S_x(\omega) = \frac{2k_B T}{\omega} \operatorname{Im} [\tilde{\alpha}(\omega)]. \quad (\text{A284})$$

[It's possible that this Appendix should also contain a derivation of the FDT.]

Problem 184: Some details about noise spectra. You may remember the formula for Johnson noise as $S_I = 4k_B T/R$, rather than the factor of 2 given above. Also, there are a few obvious differences between Eqs (A281) and (A284). Be sure you understand all these differences. The key ingredients are that all our integrals run over positive and negative frequencies, and that

while voltage is analogous to force, current is analogous to velocity, not position. Check carefully that all the details work out.

In any system at thermal equilibrium, if we apply a small force we can observe a proportionally small displacement, and this is described by a linear response function. In a mechanical system we have the function $\tilde{\alpha}(\omega)$, sometimes called a "complex compliance." In magnetic systems, the force is an applied magnetic field and the analog of position is the magnetization; the response function is called the susceptibility. Electrical systems are a bit odd because we usually discuss the current response to voltage, but we can also think about charge movements (see problem above). In all these cases, once we know the linear response function we can predict the spectral density of fluctuations in the relevant position-like variable using Eq (A284). This is called the fluctuation dissipation theorem. [Should there be an appendix with more details, and a proof? Advice welcome.]

Problem 185: Recovering equipartition. If we go to zero frequency, we have $\tilde{x} = \tilde{\alpha}(0)\tilde{F}$, but this means that $\tilde{\alpha}(0) = 1/\kappa$, where κ is the stiffness of the system. We know from the equipartition theorem that the variance in position must be related to the stiffness,

$$\frac{1}{2} \kappa \langle x^2 \rangle = \frac{1}{2} k_B T. \quad (\text{A285})$$

But we can also write the variance in position as an integral over the spectral density,

$$\langle x^2 \rangle = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} S_x(\omega). \quad (\text{A286})$$

For these equations to be consistent, we must have

$$2 \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \frac{1}{\omega} \operatorname{Im} [\tilde{\alpha}(\omega)] = \tilde{\alpha}(0), \quad (\text{A287})$$

which looks quite remarkable.

(a.) The frequency domain Eq (A282) is equivalent to

$$x(t) = \int_{-\infty}^{\infty} d\tau \alpha(\tau) F(t - \tau), \quad (\text{A288})$$

where

$$\alpha(\tau) = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} e^{-i\omega\tau} \tilde{\alpha}(\omega). \quad (\text{A289})$$

Causality means that $\alpha(\tau < 0) = 0$. What does this imply about the analytic properties of the $\tilde{\alpha}(\omega)$ in the complex ω plane?

(b.) Use your result in (a.) to verify Eq (A287).

⁹⁶ In what follows I make free use of the concepts of correlation functions, power spectra, and all that. See Appendix A.2 for a review of these ideas.

⁹⁷ Here, more than in other sections, our conventions in defining the Fourier transform are important. Be careful about the sign of i in the exponential!

Position and force, magnetization and magnetic field, charge and voltage; all of these are "thermodynamically conjugate" pairs of variables. More precisely, if we consider an ensemble in which the force is held fixed, then the derivative of the free energy with respect to the force

is the mean position, and conversely. The fluctuation dissipation theorem always refers to these pairs of variables. So, to describe fluctuations in chemical systems, we need to know the “force” that is conjugate to the concentration (or the number of molecules), and this is the chemical potential μ . To compute the response of the concentration to changes in chemical potential, we consider the diffusion equation for the concentration $c(\vec{x}, t)$ in the presence of a varying chemical potential $\mu(\vec{x}, t)$,

$$\frac{\partial c(\vec{x}, t)}{\partial t} = D \nabla \cdot \left[\nabla c(\vec{x}, t) - \frac{\nabla \mu(\vec{x})}{k_B T} c(\vec{x}, t) \right]. \quad (\text{A290})$$

Problem 186: Connecting back. Explain how Eq (A290) relates to the equation for diffusion in the presence of an external potential, Eq (240). Be sure you understand the signs.

Linearizing Equation (A290) around a mean concentration \bar{c} , we have

$$\frac{\partial c(\vec{x}, t)}{\partial t} = D \nabla^2 c(\vec{x}, t) - \frac{D \bar{c}}{k_B T} \nabla^2 \mu(\vec{x}). \quad (\text{A291})$$

We can solve by Fourier transforming in both space and time,

$$c(\vec{x}, t) = \int \frac{d^3 k}{(2\pi)^3} \int \frac{d\omega}{2\pi} e^{-i\omega t} e^{+i\vec{k} \cdot \vec{x}} \tilde{c}(\vec{k}, \omega), \quad (\text{A292})$$

where as usual \tilde{W} denotes the Fourier transform of W . If we want to identify integration with a weight W as equivalent to computing an average over the volume V , then we must have $\tilde{W}(0) = 1$, and $\tilde{W}(\vec{k})$ must decay to zero for $k \gg 1/\ell$, where ℓ is the characteristic linear dimension of the region over which we are averaging.

Equation (A299) allows us to identify the power spectrum of fluctuations in C ,

$$S_C(\omega) = \int \frac{d^3 k}{(2\pi)^3} |\tilde{W}(\vec{k})|^2 S_c(\vec{k}, \omega). \quad (\text{A300})$$

to find

$$\tilde{c}(\vec{k}, \omega) = \frac{D \bar{c}}{k_B T} \frac{k^2}{-i\omega + Dk^2} \tilde{\mu}(\vec{k}, \omega). \quad (\text{A293})$$

Thus there is a \vec{k} -dependent response function,

$$\tilde{\alpha}(\vec{k}, \omega) = \frac{D \bar{c}}{k_B T} \frac{k^2}{-i\omega + Dk^2} \quad (\text{A294})$$

from which we can use the fluctuation dissipation theorem to calculate the spatiotemporal power spectrum of concentration fluctuations,

$$S_c(\vec{k}, \omega) = \frac{2k_B T}{\omega} \text{Im} \left[\tilde{\alpha}(\vec{k}, \omega) \right] = 2\bar{c} \frac{Dk^2}{\omega^2 + (Dk^2)^2}. \quad (\text{A295})$$

Notice that the factors of $k_B T$ cancel: the fluctuations are proportional to the temperature, but the response function—the susceptibility of the concentration to changes in chemical potential—is inversely proportional to the temperature.

How does the result in Eq (A295) relate to our intuition about the \sqrt{N} rule? Let’s think about measuring the average concentration in a small volume, which corresponds to the heuristic calculation by Berg and Purcell. To do this we construct a variable

$$C(t) = \int d^3 x W(\vec{x}) c(\vec{x}, t), \quad (\text{A296})$$

where the weighting function $W(\vec{x})$ is $1/V$ inside a volume V , and zero outside. Then the correlation function of C is given by

$$\langle C(t)C(t') \rangle = \int d^3 x W(\vec{x}) \int d^3 x' W(\vec{x}') \langle c(\vec{x}, t) c(\vec{x}', t') \rangle \quad (\text{A297})$$

$$= \int d^3 x W(\vec{x}) \int d^3 x' W(\vec{x}') \int \frac{d^3 k}{(2\pi)^3} e^{i\vec{k} \cdot (\vec{x} - \vec{x}')} \int \frac{d\omega}{2\pi} e^{-i\omega(t-t')} S_c(\vec{k}, \omega) \quad (\text{A298})$$

$$= \int \frac{d\omega}{2\pi} e^{-i\omega(t-t')} \int \frac{d^3 k}{(2\pi)^3} |\tilde{W}(\vec{k})|^2 S_c(\vec{k}, \omega), \quad (\text{A299})$$

To make progress let’s assume that the region we are averaging over is spherically symmetric, so that

$$\begin{aligned} S_C(\omega) &= \int \frac{d^3 k}{(2\pi)^3} |\tilde{W}(\vec{k})|^2 S_c(\vec{k}, \omega) \\ &= \bar{c} \int \frac{d^3 k}{(2\pi)^3} |\tilde{W}(\vec{k})|^2 \frac{2Dk^2}{\omega^2 + (Dk^2)^2} \end{aligned} \quad (\text{A301})$$

$$= \bar{c} \frac{1}{(2\pi)^3} \int_0^\infty dk 4\pi k^2 |\tilde{W}(k)|^2 \frac{2Dk^2}{\omega^2 + (Dk^2)^2}. \quad (\text{A302})$$

Now if we want to compute the variance in C , we have

$$\langle(\delta C)^2\rangle \equiv \int \frac{d\omega}{2\pi} S_C(\omega) \quad (\text{A303})$$

$$= \bar{c} \frac{1}{(2\pi)^3} \int_0^\infty dk 4\pi k^2 |\tilde{W}(k)|^2 \int \frac{d\omega}{2\pi} \frac{2Dk^2}{\omega^2 + (Dk^2)^2} \quad (\text{A304})$$

$$= \bar{c} \frac{1}{2\pi^2} \int_0^\infty dk k^2 |\tilde{W}(k)|^2. \quad (\text{A305})$$

As an approximation, we can say that the effect of $|\tilde{W}(k)|^2$ is to cut the k integral off at $k \sim 2\pi/\ell$, in which case we have

$$\begin{aligned} \langle(\delta C)^2\rangle &= \bar{c} \frac{1}{2\pi^2} \int_0^\infty dk k^2 |\tilde{W}(k)|^2 \\ &\sim \bar{c} \frac{1}{2\pi^2} \int_0^{2\pi/\ell} dk k^2 \quad (\text{A306}) \\ &\sim \frac{\bar{c}}{\ell^3}. \quad (\text{A307}) \end{aligned}$$

Since $\bar{C} = \bar{c}$, we can also write this as

$$\frac{\langle(\delta C)^2\rangle}{\bar{C}^2} \sim \frac{1}{\bar{c}\ell^3}, \quad (\text{A308})$$

and we recognize $N = \bar{c}\ell^3$ as the mean number of molecules in the sampling volume. Thus, the rigorous calculation from the fluctuation dissipation theorem gives us back our intuition about the fractional variance in concentration being $1/N$.

To get the rest of the Berg–Purcell result, let's go back to Eq (A302) and finish computing the power spectrum of C , in the same approximations:

$$\begin{aligned} S_C(\omega) &= 2\bar{c} \frac{1}{(2\pi)^3} \int_0^\infty dk 4\pi k^2 |\tilde{W}(k)|^2 \frac{Dk^2}{\omega^2 + (Dk^2)^2} \\ &\sim \frac{\bar{c}}{\pi^2} \int_0^{2\pi/\ell} dk \frac{Dk^4}{\omega^2 + (Dk^2)^2}. \quad (\text{A309}) \end{aligned}$$

If ℓ is small, then the characteristic time for diffusion across the averaging volume, $\tau \sim \ell^2/D$, is also small, and hence any frequencies that are likely to be relevant for the cell's measurements of concentration are low compared with the scales on which $S_C(\omega)$ has structure. Thus we can confine our attention to the low frequency limit,

$$S_C(\omega \rightarrow 0) \sim \frac{\bar{c}}{\pi^2} \int_0^{2\pi/\ell} dk \frac{Dk^4}{(Dk^2)^2} = \frac{2\bar{c}}{\pi D\ell}. \quad (\text{A310})$$

So we see that the concentration, averaged over a sampling volume of linear dimension ℓ has white noise in time. If we average over a time τ_{avg} , then we are sensitive to a bandwidth $1/\tau_{\text{avg}}$, and we will see a variance

$$\langle(\delta C)^2\rangle_{\tau_{\text{avg}}} \sim \frac{2\bar{c}}{\pi D\ell\tau_{\text{avg}}}. \quad (\text{A311})$$

Rewriting this as a fractional standard deviation, we have

$$\frac{\delta C_{\text{rms}}}{\bar{C}} = \frac{1}{\bar{C}} \sqrt{\langle(\delta C)^2\rangle_{\tau_{\text{avg}}}} \sim \left(\frac{2}{\pi}\right)^{1/2} \frac{1}{\sqrt{D\ell\bar{c}\tau_{\text{avg}}}}, \quad (\text{A312})$$

which is (except for the trivial factor $\sqrt{2/\pi}$) exactly the Berg–Purcell result.

Problem 187: Concentration fluctuations in one dimension. Repeat the analysis we have just done, but in one dimension. Before going through a detailed calculation, you should try to anticipate the answer. We still expect (from the \sqrt{N} intuition) that $\langle(\delta C)^2\rangle \propto \bar{c}$, but since concentration has units of molecules per length in 1D, the other factors must be different. Try, for example,

$$\langle(\delta C)^2\rangle \sim \frac{\bar{c}}{(D\tau_{\text{avg}})^n \ell^m}. \quad (\text{A313})$$

How are n and m constrained by dimensional analysis? Can you argue, qualitatively, for particular values of these exponents? Finally, do the real calculation and get the analog of Eq (A311) in one dimension. Are you surprised by the role of ℓ (that is, by the value of m)? Can you explain why things come out this way?

Problem 188: Correlations seen by a moving observer. Generalize the discussion above to the case where the volume in which we measure the concentration is moving at speed v_0 in some direction. Provide a formula for the correlations across time in the observed noise. Show, in particular, that there is a correlation time $\tau_c \sim D/v_0^2$. How does this relate to the qualitative argument, discussed above, that bacteria must integrate for a minimum time $\sim D/v_0^2$ if they are to “outrun” diffusion?

So the Berg–Purcell argument certainly gives the right answer for the concentration fluctuations in a small volume. But biological systems don't actually count the molecules in a volume. Instead, the molecules bind to specific sites, and it is this binding which is detected, e.g. by activating an enzymatic reaction. The Berg–Purcell formula suggests that there is a limit to the accuracy of sensing or signaling that comes from the physics of diffusion alone, independent of these details. To see how this can happen, we need to analyze fluctuations in the binding of molecules to receptor sites, coupled to their diffusion. Let's start just with the binding events.

Consider a binding site for signaling molecules, and let the fractional occupancy of the site be n . If we do not worry about the discreteness of this one site, or about the fluctuations in concentration c of the signaling molecule, we can write a kinetic equation

$$\frac{dn(t)}{dt} = k_+c[1 - n(t)] - k_-n(t). \quad (\text{A314})$$

This describes the kinetics whereby the system comes to equilibrium, and the free energy F associated with binding is determined by detailed balance,

$$\frac{k_+c}{k_-} = \exp\left(\frac{F}{k_B T}\right). \quad (\text{A315})$$

If we imagine that thermal fluctuations can lead to small changes in the rate constants, we can linearize Eq. (A314) to obtain

$$\frac{d\delta n}{dt} = -(k_+c + k_-)\delta n + c(1 - \bar{n})\delta k_+ - \bar{n}\delta k_-. \quad (\text{A316})$$

But from Eq. (A315) we have

$$\frac{\delta k_+}{k_+} - \frac{\delta k_-}{k_-} = \frac{\delta F}{k_B T}. \quad (\text{A317})$$

Applying this constraint to Eq. (A316) we find that the individual rate constant fluctuations cancel and all that remains is the fluctuation in the thermodynamic binding energy δF :

$$\frac{d\delta n}{dt} = -(k_+c + k_-)\delta n + k_+c(1 - \bar{n})\frac{\delta F}{k_B T}. \quad (\text{A318})$$

Fourier transforming, we can solve Eq. (A318) to find the frequency dependent susceptibility of the coordinate n to its conjugate force F ,

$$\tilde{\alpha}(\omega) \equiv \frac{\delta \tilde{n}(\omega)}{\delta \tilde{F}(\omega)} = \frac{1}{k_B T} \frac{k_+c(1 - \bar{n})}{-i\omega + (k_+c + k_-)}. \quad (\text{A319})$$

Now we can compute the power spectrum of fluctuations in the occupancy n using the fluctuation dissipation theorem,

$$S_n(\omega) = \frac{2k_B T}{\omega} \text{Im} \left[\frac{\delta \tilde{n}(\omega)}{\delta \tilde{F}(\omega)} \right] \quad (\text{A320})$$

$$= \frac{2k_+c(1 - \bar{n})}{\omega^2 + (k_+c + k_-)^2}. \quad (\text{A321})$$

It is convenient to rewrite this as

$$S_n(\omega) = \langle (\delta n)^2 \rangle \frac{2\tau_c}{1 + (\omega\tau_c)^2}, \quad (\text{A322})$$

where the total variance is

$$\langle (\delta n)^2 \rangle = \int \frac{d\omega}{2\pi} S_n(\omega) = k_B T \left. \frac{\delta \tilde{n}(\omega)}{\delta \tilde{F}(\omega)} \right|_{\omega=0} \quad (\text{A323})$$

$$= \frac{k_+c(1 - \bar{n})}{k_+c + k_-} \quad (\text{A324})$$

$$= \bar{n}(1 - \bar{n}), \quad (\text{A325})$$

and the correlation time is given by

$$\tau_c = \frac{1}{k_+c + k_-}. \quad (\text{A326})$$

To make sense out of these results, remember what happens if we flip a coin that is biased to produce heads a fraction f of the time. On each trial we count either one or zero heads, so the mean count is f and the mean square count is also f ; the variance is $f(1 - f)$, exactly as in Eq (A325): when we check the occupancy of the

receptor, the outcome is determined by the equivalent of flipping a biased coin, where the bias is determined by the Boltzmann distribution.

The Lorentzian form of the power spectrum in Eq (A322) is equivalent to an exponential decay of correlations,

$$\langle \delta n(t)\delta n(t') \rangle = \int \frac{d\omega}{2\pi} e^{-i\omega(t-t')} S_n(\omega) \quad (\text{A327})$$

$$= \langle (\delta n)^2 \rangle \int \frac{d\omega}{2\pi} e^{-i\omega(t-t')} \frac{2\tau_c}{1 + (\omega\tau_c)^2} \quad (\text{A328})$$

$$= \langle (\delta n)^2 \rangle e^{-|t-t'|/\tau_c}. \quad (\text{A329})$$

The exponential decay of correlations is what we expect when the transitions between the available states have no memory. To be precise about this, if we imagine that a system is in one state at time $t = 0$, and there is some constant probability per unit time k of transitions out of this state (with, in the simplest case, no returns to the initial state), then the probability $p(t)$ of still being in the initial state at time t must obey

$$\frac{dp(t)}{dt} = -kp(t), \quad (\text{A330})$$

and hence $p(t) = e^{-kt}$. This intuition about the connection of exponential decays to the lack of memory is very general, and should remind you of the exponential distribution of times between transitions in the calculation of chemical reaction rates (Section II.A), and of the exponential distribution of times between events in a Poisson process (see Appendix A.1). In the present context, the exponential decay of correlations tell us that the spontaneous transitions between the occupied and unoccupied states of the receptor occur with constant probability per unit time, or as Markovian jumps. The jumping rates are just the rates k_+ and k_- , which means that when we write chemical kinetic models for a whole ensemble of molecules, we also can interpret these as Markov models for transitions among the states of individual molecules in the ensemble.

It is interesting that we recover the results for Markovian jumping between two states without making this microscopic model explicit. All we assume is the macroscopic kinetics and that the system is in thermal equilibrium so that we can apply the fluctuation dissipation theorem. In principle many different microscopic models can describe the molecular phenomena that are at the basis of some observed macroscopic behavior, and we know that many aspects of behavior in thermal equilibrium are independent of these details. The statistics of fluctuations in a chemical kinetic system are an example of this, at least near equilibrium.

The good news, then, is that fluctuations in receptor occupancy are an inevitable consequence of the *macroscopic*, average behavior of receptor–ligand interactions, independent of hypothesis about molecular details. The

bad news is that the form of the results doesn't seem very related to the ideas of Berg and Purcell about the precision of concentration measurements. To make these connections clear we need to couple the dynamics of receptor occupancy to the diffusion of the ligand.

When the concentration is allowed to fluctuate we write

$$\frac{dn(t)}{dt} = k_+c(\vec{x}_0, t)[1 - n(t)] - k_-n(t), \quad (\text{A331})$$

where the receptor is located at \vec{x}_0 , and

$$\frac{\partial c(\vec{x}, t)}{\partial t} = D\nabla^2 c(\vec{x}, t) - \delta(\vec{x} - \vec{x}_0) \frac{dn(t)}{dt}. \quad (\text{A332})$$

The first equation is as before, but with notation to remind us that the concentration c is dynamic. The second equation states that the ligand diffuses with diffusion constant D , and when the receptor located at \vec{x}_0 increases its occupancy it removes exactly one molecule from solution at that point.

Problem 189: Coupling diffusion and binding. In this problem you'll fill the details needed for the analysis of Eq's (A331) and (A332).

(a.) Begin by noticing that Eq (A332) is linear, so you should be able to solve it exactly. Use Fourier transforms, both in space and time, and then transform back to give a formal expression for

$$\tilde{c}(\vec{x}, \omega) = \int dt e^{+i\omega t} c(\vec{x}, t). \quad (\text{A333})$$

(b.) Linearize Eq (A331), in the same way that we did in the preceding derivation, leading from Eq (A314) to (A319). Along

$$\frac{\delta \tilde{n}(\omega)}{\delta \tilde{F}(\omega)} = \frac{k_+ \bar{c}(1 - \bar{n})}{k_B T} \left[-i\omega \left(1 + \frac{k_+(1 - \bar{n})}{2\pi D a} \right) + (k_+ \bar{c} + k_-) \right]^{-1}, \quad (\text{A337})$$

where \bar{c} is the mean concentration. Applying the fluctuation-dissipation theorem once again we find the spectral density of occupancy fluctuations,

$$S_n(\omega) \approx 2k_+ \bar{c}(1 - \bar{n}) \frac{1 + \Sigma(0)}{\omega^2(1 + \Sigma(0))^2 + (k_+ \bar{c} + k_-)^2}. \quad (\text{A338})$$

The total variance in occupancy is unchanged, since this is an equilibrium property of the system, while coupling to concentration fluctuations serves only to change the kinetics.

Problem 190: Reading off the results. You should be able to verify the statements in the last sentence without detailed cal-

the way you will need an expression for $\tilde{c}(\vec{x}_0, \omega)$, which you can take from (a.). When the dust settles, you should find Eq's (A334, A335)

Following the same steps as above, we find the linear response function

$$\frac{\delta \tilde{n}(\omega)}{\delta \tilde{F}(\omega)} = \frac{k_+ c(1 - \bar{n})}{k_B T} \frac{1}{-i\omega[1 + \Sigma(\omega)] + (k_+ \bar{c} + k_-)} \quad (\text{A334})$$

$$\Sigma(\omega) = k_+(1 - \bar{n}) \int \frac{d^3 k}{(2\pi)^3} \frac{1}{-i\omega + Dk^2} \quad (\text{A335})$$

The "self-energy" $\Sigma(\omega)$ is ultraviolet divergent, which can be traced to the delta function in Eq (A332); we have assumed that the receptor is infinitely small. A more realistic treatment would give the receptor a finite size, which is equivalent to cutting off the k integrals at some (large) $\Lambda \sim \pi/a$, with a the linear dimension of the receptor. If we imagine mechanisms which read out the receptor occupancy average over a time τ long compared to the correlation time τ_c of the noise, then the relevant quantity is the low frequency limit of the noise spectrum. Hence,

$$\Sigma(\omega \ll D/a^2) \approx \Sigma(0) = \frac{k_+(1 - \bar{n})}{2\pi D a}, \quad (\text{A336})$$

and

ulation. Explain how to "read" Eq (A338) and identify the total variance and correlation time.

Coupling to concentration fluctuations does serve to renormalize the correlation time of the noise,

$$\tau_c \rightarrow \tau_c[1 + \Sigma(0)]. \quad (\text{A339})$$

The new τ_c can be written as

$$\tau_c = \frac{1 - \bar{n}}{k_-} + \frac{\bar{n}(1 - \bar{n})}{2\pi D a \bar{c}}, \quad (\text{A340})$$

so there is a lower bound on τ_c , independent of the kinetic parameters k_{\pm} ,

$$\tau_c > \frac{\bar{n}(1-\bar{n})}{2\pi Da\bar{c}}. \quad (\text{A341})$$

As discussed previously, the relevant quantity is the low frequency limit of the noise spectrum,

$$S_n(\omega = 0) = 2k_+\bar{c}(1-\bar{n}) \cdot \frac{1 + \Sigma(0)}{(k_+\bar{c} + k_-)^2} \quad (\text{A342})$$

$$= \frac{2\bar{n}(1-\bar{n})}{k_+\bar{c} + k_-} + \frac{[\bar{n}(1-\bar{n})]^2}{\pi Da\bar{c}}. \quad (\text{A343})$$

If we average for a time τ , then the root-mean-square error in our estimate of n will be

$$\delta n_{\text{rms}} = \sqrt{S_n(0)} \cdot \frac{1}{\tau}, \quad (\text{A344})$$

and we see that this noise level has a minimum value independent of the kinetic parameters k_{\pm} ,

$$\delta n_{\text{rms}} > \frac{\bar{n}(1-\bar{n})}{\sqrt{\pi Da\bar{c}\tau}}. \quad (\text{A345})$$

To relate these results back to the discussion by Berg and Purcell, we note that the $\omega = 0$ response of the mean occupancy to changes in concentration can be written as

$$\frac{d\bar{n}}{d \ln c} = \bar{n}(1-\bar{n}). \quad (\text{A346})$$

Thus, the fluctuations in n are equivalent to fluctuations in c :

$$\frac{\delta c_{\text{eff}}}{\bar{c}} = (\delta \ln c)_{\text{eff}} = \delta n_{\text{rms}} \left(\frac{d\bar{n}}{d \ln c} \right)^{-1} = \frac{1}{\sqrt{\pi Da\bar{c}\tau}}. \quad (\text{A347})$$

Except for the factor of $\sqrt{\pi}$, this is the Berg–Purcell result once again.

A startling feature of the Berg–Purcell argument is that (it seems) it can be used both when a is the size of a single receptor molecule and when a is the size of the entire bacterium. Naively, we might expect that if there are N receptors on the surface of the cell, then the signal-to-noise ratio for concentration measurements should be N times better, and correspondingly the threshold for reliable detection should be \sqrt{N} times smaller,

$$\frac{\delta c_{\text{eff}}}{\bar{c}} \sim \frac{1}{\sqrt{DNA\bar{c}\tau}}. \quad (\text{A348})$$

On the other hand, if we use the Berg–Purcell limit and take the linear dimensions of the detector to be the radius R of the bacterium, we should obtain

$$\frac{\delta c_{\text{eff}}}{\bar{c}} \sim \frac{1}{\sqrt{DR\bar{c}\tau}}. \quad (\text{A349})$$

What is going on? Does something special happen when $N \sim R/a$, so there is a crossover between the two results?

If we imagine a very large cell, and place $N = 2$ two receptors on opposite sides of the cell surface, it is hard to imagine that there is anything wrong with the argument leading to Eq (A348). More generally, if the receptors are far apart, it is very plausible that they report independent measurements of the concentration, and so Eq (A348) should be correct. On the other hand, if we imagine bringing two receptors closer and closer together, at some point they will start to interact—a molecule released from one receptor can diffuse over and bind to the other receptor—and this interaction might lead to correlations in the noise, and a break down of the simple \sqrt{N} improvement in the threshold for reliable detection.




FIG. 174 Correlated Brownian motion, from Meiners & Quake (1999). At left, a schematic of the experiment. The laser beams from the bottom of the figure create two optical traps, which hold the microspheres in approximately harmonic potential wells. The optics at the top allow for measurements of the spheres’ positions with nanometer precision. At right, measurements of the auto- and cross-correlations of the spheres’ positions; different curves for the cross-correlation correspond to different mean separations of the particles, which is expected to modulate the coupling between them through the fluid.

To understand how diffusive interactions lead to correlations among receptors, it is useful to think about a simpler problem. Suppose that we have two balls in a fluid. If they are very far apart, each one experiences a drag force and undergoes Brownian motion, and the Brownian fluctuations in the position of each ball are independent of those in the other. If we bring the two balls together, however, we know that they can influence each other through the fluid: If one ball moves at velocity v_1 it not only experiences a drag force $-\gamma v_1$, it also applies a “coupling” force $\gamma_c(v_1 - v_2)$ to the other ball (which may be moving at velocity v_2 ; clearly if $v_1 = v_2$ there should be no coupling force). If the balls are close enough that γ_c is significant, then in fact the Brownian motions of the two balls become correlated. This correlation can be

derived from the fluctuation–dissipation theorem, and it also makes intuitive sense since a random Brownian step of one object applies a force to the other. We can also see this effect experimentally, as in Fig 174.

Problem 191: Correlated Brownian motion. To make the situation in the previous paragraph precise, consider the case where the particles are bound by springs (so they can’t diffuse away from each other and reduce the coupling). Then, in the overdamped case, the equations of motion are

$$\gamma \frac{dx_1}{dt} = -\kappa x_1 - \gamma_c \left(\frac{dx_1}{dt} - \frac{dx_2}{dt} \right) + F_1(t) \quad (\text{A350})$$

$$\gamma \frac{dx_2}{dt} = -\kappa x_2 - \gamma_c \left(\frac{dx_2}{dt} - \frac{dx_1}{dt} \right) + F_1(t) \quad (\text{A351})$$

where κ is the stiffness of the springs (assumed identical, for simplicity), and $F_i(t)$ is an external force applied to each particle i .

(a.) Derive the linear response function matrix, $\tilde{\alpha}_{ij}(\omega)$ such that

$$\tilde{x}_i(\omega) = \sum_j \tilde{\alpha}_{ij}(\omega) \tilde{F}_j(\omega). \quad (\text{A352})$$

(b.) The generalization of the fluctuation dissipation theorem to many degrees of freedom states that the “cross–spectrum” of variables i and j , defined by

$$\langle x_i(t)x_j(t') \rangle = \int \frac{d\omega}{2\pi} e^{-i\omega(t-t')} S_{ij}(\omega), \quad (\text{A353})$$

is given by

$$S_{ij}(\omega) = \frac{2k_B T}{\omega} \text{Im} [\tilde{\alpha}_{ij}(\omega)]. \quad (\text{A354})$$

Use this to derive the cross–spectrum of the position fluctuations for the two particles.

(c.) Despite the viscous coupling, the potential energy is just the sum of contributions from the two particles. From the Boltzmann distribution, then, the positions should be independent variables. Use your results in (b) to show that $\langle x_i x_j \rangle = \delta_{ij} k_B T / \kappa$. Notice that this corresponds to the *instantaneous* positions of the particles, as we would measure by taking a snapshot (with a fast camera).

(d.) Suppose that instead of taking snapshots of the positions, we average (as in the discussion above) for a long time, so what is relevant is the low frequency limit of the power spectra. Show that now the correlations are nonzero, and give an explicit formula for the covariance matrix of fluctuations in the temporally averaged positions.

If we imagine that positions of the Brownian particles are like receptor occupancies, and an applied force on all the particles is like a change in concentration of the relevant ligand, then diffusion of the ligand serves the same coupling effect as the viscosity of the fluid and will generate correlations among the occupancy fluctuations of nearby receptors. These correlations mean that using the positions or velocities of N Brownian particles to infer the applied force is *not* \sqrt{N} more accurate than using one particle, and similarly using N receptors will not generate a concentration measurement that is \sqrt{N} times more accurate than is obtained with one receptor.

If we have N receptors, each of size a arrayed on a structure of linear dimension R such as a ring or a sphere, then as N becomes large the receptors are coming closer and closer together, and we expect that correlations become stronger. If we have two detectors making measurements with noise that becomes more and more strongly correlated, at some point they start to act like one big detector. If we work through the details of the calculations for the case of multiple receptors,⁹⁸ indeed we find that as N become large, the correlations among the different receptors become limiting, and the threshold for reliable detection approaches Eq (A349): the $N \rightarrow \infty$ receptors packed into a structure with linear dimension R acts like one receptor of size R . If we go back to the intuitive Berg–Purcell argument about counting molecules in a volume and getting a fresh count each time the volume clears from diffusion, what this means is that packing many receptor sites into a region of size R eventually means that we get to count the molecules in a volume $\sim R^3$. There are geometrical factors for different spatial arrangements of the receptors, but like the $\sqrt{\pi}$ in Eq (A347) these aren’t a big deal.

Almost all of PhD students in physics have seen some cases of the fluctuation dissipation theorem, somewhere in their statistical mechanics courses. Whether you have seen the general formulation depends a bit on who taught the course, and how far you went. As usual, an excellent discussion can be found in Landau & Lifshitz (1977). A later volume in the Landau and Lifshitz series (Lifshitz & Pitaevskii 1980) provides a clear discussion of concentration fluctuations, in Section 89. Many people find the idea of correlations between Brownian particles to be surprising, so it’s worth looking at real experiments that measure these correlations (Meiners & Quake 1999). [\[should add refs to measurements on concentration fluctuations—Feher and Weissman?\]](#)

Landau & Lifshitz 1977: *Statistical Physics*. LD Landau & EM Lifshitz (Pergamon, Oxford, 1977).

Lifshitz & Pitaevskii 1980: *Statistical Physics, Part 2*. EM Lifshitz & LP Pitaevskii (Pergamon, Oxford, 1980).

Meiners & Quake 1999: Direct measurement of hydrodynamic cross correlations between two particles in an external potential. JC Meiners & SR Quake, *Phys Rev Lett* **82**, 2211–2214 (1999).

The idea that fluctuations in certain chemical systems could be described using the fluctuation dissipation theorem must have occurred to many people, and I remember discussing it long ago (Bialek 1987). The emergence of experiments on noise in the control of gene expression made it more interesting to get everything straight, so my colleagues and I did this in a series of papers (Bialek & Setayeshgar 2005, 2008; Tkačik & Bialek 2009).

Bialek 1987: Physical limits to sensation and perception. W Bialek, *Ann Rev Biophys Biophys Chem* **16**, 455–478 (1987).

⁹⁸ See the references at the end of this section for details.

- Bialek & Setayeshgar 2005:** Physical limits to biochemical signaling. W Bialek & S Setayeshgar, *Proc Nat'l Acad Sci (USA)* **102**, 10040–10045 (2005).
- Bialek & Setayeshgar 2008:** Cooperativity, sensitivity and noise in biochemical signaling. W Bialek & S Setayeshgar, *Phys Rev Lett* **100**, 258101 (2008).
- Tkačik & Bialek 2009:** Diffusion, dimensionality and noise in transcriptional regulation. G Tkačik & W Bialek, *Phys Rev E* **79**, 051901 (2009).

7. Dimensionality reduction

I am leaving this unwritten for now, in the interests of getting something readable out more quickly. I think it is straightforward to write. To give a sense of what goes here, I have started to compile the references. Evidently papers by my colleagues and myself are over-represented; of course a full account will look at a broader literature.

Problem 192: Analysis of a sensory neuron. [Get a big data set from Rob on H1, and use it to take the students through reverse correlation, spike triggered covariance and (maybe) maximally informative dimensions. Have repeats so one can compare reduced models with the real information per spike.]

Problem 193: Analysis of DNA sequences. [Get data from Justin and take the students through a small version of the problem (maybe the RNAP site).]

- d'Avella & Bizzi 1998:** Low dimensionality of surpaspinally induced force fields. A d'Avella & E Bizzi, *Proc Nat'l Acad Sci (USA)* **95**, 7711–7714 (1998).
- Bialek & de Ruyter van Steveninck 2005:** Features and dimensions: Motion estimation in fly vision. W Bialek & R de Ruyter van Steveninck, arXiv:q-bio/0505003 (2005).
- de Boer & Kuyper 1968:** Triggered correlation. E de Boer & P Kuyper, *IEEE Trans Biomed Eng* **15**, 169–179 (1968).
- Chigirev & Bialek 2004:** Optimal manifold representation of data: An information theoretic perspective. DV Chigirev & W Bialek, in *Advances in Neural Information Processing 16*, S Thrun, L Saul & B Schölkopf, eds, pp 161–168 (MIT Press, Cambridge, 2004).
- Fairhall et al 2006:** Selectivity for multiple stimulus features in retinal ganglion cells. AL Fairhall, CA Burlingame, R Narasimhan, RA Harris, JL Puchalla & MJ Berry II, *J Neurophysiol* **96**, 2724–2738 (2006).

- Kinney et al 2007:** Precise physical models of protein–DNA interaction from high-throughput data. JB Kinney, G Tkačik & CG Callan Jr, *Proc Nat'l Acad Sci (USA)* **104**, 501–506 (2007).

- Kinney et al 2010:** Using deep sequencing to characterize the biophysical mechanism of a transcriptional regulatory sequence. JB Kinney, A Murugan, CG Callan Jr & EC Cox, *Proc Nat'l Acad Sci (USA)* **107**, 9158–9163 (2010).

- Osborne et al 2005:** A sensory source for motor variation. LC Osborne, SG Lisberger & W Bialek, *Nature* **437**, 412–416 (2005).

- Rust et al 2005:** Spatiotemporal elements of macaque V1 receptive fields. NC Rust, O Schwartz, JA Movshon & EP Simoncelli, *Neuron* **46**, 945–956 (2005).

- de Ruyter van Steveninck & Bialek 1988:** Real-time performance of a movement sensitive neuron in the blowfly visual system: Coding and information transfer in short spike sequences. R de Ruyter van Steveninck & W Bialek, *Proc R. Soc London Ser. B* **234**, 379–414 (1988).

- Sanger 2000:** Human arm movements described by a low-dimensional superposition of principal components. TD Sanger, *J Neurosci* **20**, 1066–1072 (2000).

- Sharpee et al 2004:** Analyzing neural responses to natural signals: Maximally informative dimensions. T Sharpee, NC Rust & W Bialek, *Neural Comp* **16**, 223–250 (2004); arXiv:physics/0212110 (2002).

- Stephens et al 2008:** Dimensionality and dynamics in the behavior of *C. elegans*. GJ Stephens, B Johnson–Kerner, W Bialek & WS Ryu, *PLoS Comp Bio* **4**, e1000028 (2008); arXiv:0705.1548 [q-bio.OT] (2007).

- Stephens et al 2011:** Searching for simplicity in the analysis of neurons and behavior. GJ Stephens, LC Osborne & W Bialek, *Proc Nat'l Acad Sci (USA)* in press (2011); arXiv.org:1012.3896 [q-bio.NC] (2010).

8. Maximum entropy

This section is a bit old. It needs to be revised in light of what happens in Sections III.A and III.D. Be sure that we go into RESULTS on these methods, as promised for neurons, at least.

The problem of finding the maximum entropy given some constraint again is familiar from statistical mechanics: the Boltzmann distribution is the distribution that has the largest possible entropy given the mean energy. More generally, imagine that we have knowledge not of the whole probability distribution $P(D)$ but only of some expectation values,

$$\langle f_i \rangle = \sum_D P(D) f_i(D), \quad (\text{A355})$$

where we allow that there may be several expectation values known ($i = 1, 2, \dots, K$). Actually there is one more expectation value that we always know, and this is that

the average value of one is one; the distribution is normalized:

$$\langle f_0 \rangle = \sum_D P(D) = 1. \quad (\text{A356})$$

Given the set of numbers $\{\langle f_0 \rangle, \langle f_1 \rangle, \dots, \langle f_K \rangle\}$ as constraints on the probability distribution $P(D)$, we would like to know the largest possible value for the entropy,

and we would like to find explicitly the distribution that provides this maximum.

The problem of maximizing a quantity subject to constraints is formulated using Lagrange multipliers. In this case, we want to maximize $S = -\sum P(D) \log_2 P(D)$, so we introduce a function \tilde{S} , with one Lagrange multiplier $\tilde{\lambda}_i$ for each constraint:

$$\tilde{S}[P(D)] = -\sum_D P(D) \log_2 P(D) - \sum_{i=0}^K \tilde{\lambda}_i \langle f_i \rangle \quad (\text{A357})$$

$$= -\frac{1}{\ln 2} \sum_D P(D) \ln P(D) - \sum_{i=0}^K \lambda_i \sum_D P(D) f_i(D). \quad (\text{A358})$$

Our problem, then, is to find the maximum of the function \tilde{S} , but this is easy because the probability for each value of D appears independently. As usual, we differentiate and set the result to zero:

$$0 = \frac{\partial \tilde{S}}{\partial P(D)} = -\frac{1}{\ln 2} [\ln P(D) + 1] - \sum_{i=0}^K \tilde{\lambda}_i f_i(D). \quad (\text{A359})$$

Rearranging, we have

$$\ln P(D) = -1 - \sum_{i=0}^K (\ln 2) \tilde{\lambda}_i f_i(D) \quad (\text{A360})$$

$$P(D) = \frac{1}{Z} \exp \left[-\sum_{i=0}^K \lambda_i f_i(D) \right], \quad (\text{A361})$$

where $\lambda_i = (\ln 2) \tilde{\lambda}_i$, and $Z = \exp(1 + \lambda_0)$ is a normalization constant. Notice that this gives us the *form* of the maximum entropy distribution, but we still have to adjust the constants $\{\lambda_i\}$ so that the distribution $P(D)$ predicts the measured values of the expectation values in Eq (A355).

There are several things worth saying about maximum entropy distributions. First, we recall that if the value of D indexes the states n of a physical system, and we know only the expectation value of the energy,

$$\langle E \rangle = \sum_n P_n E_n, \quad (\text{A362})$$

then the maximum entropy distribution is

$$P_n = \frac{1}{Z} \exp(-\lambda E_n), \quad (\text{A363})$$

which is the Boltzmann distribution (as promised). In this case the Lagrange multiplier λ has physical meaning—it is the inverse temperature. Further, the function \tilde{S} that we introduced for convenience is the difference between the entropy and λ times the energy; if

we divide through by λ and flip the sign, then we have the energy minus the temperature times the entropy, or the free energy. Thus the distribution which maximizes entropy at fixed average energy is also the distribution which minimizes the free energy.

If we are looking at a magnetic system, for example, and we know not just the average energy but also the average magnetization, then a new term appears in the exponential of the probability distribution, and we can interpret this term as the magnetic field multiplied by the magnetization. More generally, for every order parameter which we assume is known, the probability distribution acquires a term that adds to the energy and can be thought of as a product of the order parameter with its conjugate force. Again, all these remarks should be familiar from a statistical mechanics course.

Consider the situation in which the data D are real numbers x . Suppose that we know the mean value of x and its variance. This is equivalent to knowledge of two expectation values,

$$\bar{f}_1 = \langle x \rangle = \int dx P(x) x, \quad \text{and} \quad (\text{A364})$$

$$\bar{f}_2 = \langle x^2 \rangle = \int dx P(x) x^2, \quad (\text{A365})$$

so we have $f_1(x) = x$ and $f_2(x) = x^2$. Thus, from Eq. (A361), the maximum entropy distribution is of the form

$$P(x) = \frac{1}{Z} \exp(-\lambda_1 x - \lambda_2 x^2). \quad (\text{A366})$$

This is a funny way of writing a more familiar object. If we identify the parameters $\lambda_2 = 1/(2\sigma^2)$ and $\lambda_1 = -\langle x \rangle/\sigma^2$, then we can rewrite the maximum entropy distribution as the usual Gaussian,

$$P(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2\sigma^2}(x - \langle x \rangle)^2\right]. \quad (\text{A367})$$

We recall that Gaussian distributions usually arise through the central limit theorem: if the random variables of interest can be thought of as sums of many independent events, then the distributions of the observable variables converge to Gaussians. This provides us with a ‘mechanistic’ or reductionist view of why Gaussians are so important. A very different view comes from information theory: if all we know about a variable is the mean and the variance, then the Gaussian distribution is the maximum entropy distribution consistent with this knowledge. Since the entropy measures (returning to our physical intuition) the randomness or disorder of the system, the Gaussian distribution describes the ‘most random’ or ‘least structured’ distribution that can generate the known mean and variance.

Problem 194: Less than maximum entropy. Many natural signals are strongly nonGaussian. In particular exponential (or nearly exponential) distribution are common in studies on the statistics of natural images and natural sounds. With the same mean (which you can call zero) and variance, what is the difference in entropy between the exponential [$P(x) \propto \exp(-\lambda|x|)$] and Gaussian distributions? If we imagine that this difference is relevant to every pixel (or to every Fourier component) in an image, is this significant compared to the 8 bits/pixel of a standard digital image? What if $P(x) \propto \exp(-\lambda|x|^\mu)$, with $\mu < 1$?

[maybe we should put the start of networks here?]

[maybe this should be connections, more generally (including what we have to say about counting), and that would leave a section to address the experimental situation more specifically?]

Probability distributions that have the maximum entropy form of Eq. (A361) are special not only because of their connection to statistical mechanics, but because they form what the statisticians call an ‘exponential family,’ which seems like an obvious name. The important point is that exponential families of distributions are (almost) unique in having sufficient statistics. To understand what this means, consider the following problem: we observe a set of samples D_1, D_2, \dots, D_N , each of which is drawn independently and at random from a distribution $P(D|\{\lambda_i\})$. Assume that we know the form of this distribution but not the values of the parameters $\{\lambda_i\}$. How can we estimate these parameters from the set of observations $\{D_n\}$? Notice that our data set $\{D_n\}$ consists of N numbers, and N can be very large; on the other hand there typically are a small number $K \ll N$ of parameters λ_i that we want to estimate. Even in this limit, no finite amount of data will tell us the exact values of the parameters, and so we need a probabilistic formulation: we want to compute the distribution of parameters given the data, $P(\{\lambda_i\}|\{D_n\})$. We do this using Bayes’ rule,

$$P(\{\lambda_i\}|\{D_n\}) = \frac{1}{P(\{D_n\})} \cdot P(\{D_n\}|\{\lambda_i\})P(\{\lambda_i\}), \quad (\text{A368})$$

where $P(\{\lambda_i\})$ is the distribution from which the parameter values themselves are drawn. Then since each datum D_n is drawn independently, we have

$$P(\{D_n\}|\{\lambda_i\}) = \prod_{n=1}^N P(D_n|\{\lambda_i\}). \quad (\text{A369})$$

For probability distributions of the maximum entropy form we can proceed further, using Eq. (A361):

$$\begin{aligned} P(\{\lambda_i\}|\{D_n\}) &= \frac{1}{P(\{D_n\})} \cdot P(\{D_n\}|\{\lambda_i\})P(\{\lambda_i\}) \\ &= \frac{P(\{\lambda_i\})}{P(\{D_n\})} \prod_{n=1}^N P(D_n|\{\lambda_i\}) \end{aligned} \quad (\text{A370})$$

$$= \frac{P(\{\lambda_i\})}{Z^N P(\{D_n\})} \prod_{n=1}^N \exp\left[-\sum_{i=1}^K \lambda_i f_i(D_n)\right] \quad (\text{A371})$$

$$= \frac{P(\{\lambda_i\})}{Z^N P(\{D_n\})} \exp\left[-N \sum_{i=1}^K \lambda_i \frac{1}{N} \sum_{n=1}^N f_i(D_n)\right]. \quad (\text{A372})$$

We see that *all* of the information that the data points

$\{D_n\}$ can give about the parameters λ_i is contained in

the average values of the functions f_i over the data set, or the ‘empirical means’ \bar{f}_i ,

$$\bar{f}_i = \frac{1}{N} \sum_{n=1}^N f_i(D_n). \quad (\text{A373})$$

More precisely, the distribution of possible parameter values consistent with the data depends not on all details of the data, but rather only on the empirical means $\{\bar{f}_i\}$,

$$P(\{\lambda_i\} | D_1, D_2, \dots, D_N) = P(\{\lambda_i\} | \{\bar{f}_i\}), \quad (\text{A374})$$

and a consequence of this is the information theoretic statement

$$I(D_1, D_2, \dots, D_N \rightarrow \{\lambda_i\}) = I(\{\bar{f}_i\} \rightarrow \{\lambda_i\}). \quad (\text{A375})$$

This situation is described by saying that the reduced set of variables $\{\bar{f}_i\}$ constitute *sufficient statistics* for learning the distribution. Thus, for distributions of this form, the problem of compressing N data points into $K \ll N$ variables that are relevant for parameter estimation can be solved explicitly: if we keep track of the running averages \bar{f}_i we can compress our data as we go along, and we are guaranteed that we will never need to go back and examine the data in more detail. A clear example is that if we know data are drawn from a Gaussian distribution, running estimates of the mean and variance contain all the information available about the underlying parameter values.

The Gaussian example makes it seem that the concept of sufficient statistics is trivial: of course if we know that data are chosen from a Gaussian distribution, then to identify the distribution all we need to do is to keep track of two moments. Far from trivial, this situation is quite unusual. Most of the distributions that we might write down do not have this property—even if they are described by a finite number of parameters, we cannot guarantee that a comparably small set of empirical expectation values captures all the information about the parameter values. If we insist further that the sufficient statistics be additive and permutation symmetric, then it is a theorem that *only* exponential families have sufficient statistics.

[say more about this!]

[where do we put connection of matching expectation values to maximum likelihood?]

The generic problem of information processing, by the brain or by a machine, is that we are faced with a huge quantity of data and must extract those pieces that are of interest to us. The idea of sufficient statistics is intriguing in part because it provides an example where this problem of ‘extracting interesting information’ can be solved completely: if the points D_1, D_2, \dots, D_N are chosen independently and at random from some distribution, the only thing which could possibly be ‘interesting’ is the structure of the distribution itself (everything else

is random, by construction), this structure is described by a finite number of parameters, and there is an explicit algorithm for compressing the N data points $\{D_n\}$ into K numbers that preserve all of the interesting information. The crucial point is that this procedure cannot exist in general, but only for certain classes of probability distributions. This is an introduction to the idea some kinds of structure in data are learnable from random examples, while other structures are not.

Consider the (Boltzmann) probability distribution for the states of a system in thermal equilibrium. If we expand the Hamiltonian as a sum of terms (operators) then the family of possible probability distributions is an exponential family in which the coupling constants for each operator are the parameters analogous to the λ_i above. In principle there could be an infinite number of these operators, but for a given class of systems we usually find that only a finite set are “relevant” in the renormalization group sense: if we write an effective Hamiltonian for coarse grained degrees of freedom, then only a finite number of terms will survive the coarse graining procedure. If we have only a finite number of terms in the Hamiltonian, then the family of Boltzmann distributions has sufficient statistics, which are just the expectation values of the relevant operators. This means that the expectation values of the relevant operators carry all the information that the (coarse grained) configuration of the system can provide about the coupling constants, which in turn is information about the identity or microscopic structure of the system. Thus the statement that there are only a finite number of relevant operators is also the statement that a finite number of expectation values carries all the information about the microscopic dynamics. The ‘if’ part of this statement is obvious: if there are only a finite number of relevant operators, then the expectation values of these operators carry all the information about the identity of the system. The statisticians, through the theorem about the uniqueness of exponential families, give us the ‘only if’: a finite number of expectation values (or correlation functions) can provide all the information about the system *only if* the effective Hamiltonian has a finite number of relevant operators. I suspect that there is more to say along these lines.

An important example of the maximum entropy idea arises when the data D are generated by counting. Then the relevant variable is an integer $n = 0, 1, 2, \dots$, and it is natural to imagine that what we know is the mean count $\langle n \rangle$. One way this problem can arise is that we are trying to communicate and are restricted to sending discrete or quantized units. An obvious case is in optical communication, where the quanta are photons. In the brain, quantization abounds: most neurons do not generate continuous analog voltages but rather communicate with one another through stereotyped pulses or spikes, and even if the voltages vary continuously transmission across a synapse involves the release of a chem-

ical transmitter which is packaged into discrete vesicles. It can be relatively easy to measure the mean rate at which discrete events are counted, and we might want to know what bounds this mean rate places on the ability of the cells to convey information. Alternatively, there is an energetic cost associated with these discrete events—generating the electrical currents that underlie the spike, constructing and filling the vesicles, ... —and we might want to characterize the mechanisms by their cost per bit rather than their cost per event [Laughlin et al 1998, Sarpeshkar 1998].

If we know the mean count, there is (as for the Boltzmann distribution) only one function $f_1(n) = n$ that can appear in the exponential of the distribution, so that

$$P(n) = \frac{1}{Z} \exp(-\lambda n). \quad (\text{A376})$$

Of course we have to choose the Lagrange multiplier to fix the mean count, and it turns out that $\lambda = \ln(1 + 1/\langle n \rangle)$ [do the calculation of λ !]; further we can find the entropy

$$S_{\max}(\text{counting}) = \log_2(1 + \langle n \rangle) + \langle n \rangle \log_2(1 + 1/\langle n \rangle). \quad (\text{A377})$$

The information conveyed by counting something can never exceed the entropy of the distribution of counts, and if we know the mean count then the entropy can never exceed the bound in Eq. (A377). Thus, if we have a system in which information is conveyed by counting discrete events, the simple fact that we count only a limited number of events (on average) sets a bound on how much information can be transmitted. We will see that real neurons and synapses approach this fundamental limit.

One might suppose that if information is coded in the counting of discrete events, then each event carries a certain amount of information. In fact this is not quite right.

In particular, if we count a large number of events then the maximum counting entropy becomes

$$S_{\max}(\text{counting}; \langle n \rangle \rightarrow \infty) \sim \log_2(\langle n \rangle e), \quad (\text{A378})$$

and so we are guaranteed that the entropy (and hence the information) per event goes to zero, although the approach is slow. On the other hand, if events are very rare, so that the mean count is much less than one, we find the maximum entropy per event

$$\frac{1}{\langle n \rangle} S_{\max}(\text{counting}; \langle n \rangle \ll 1) \sim \log_2\left(\frac{e}{\langle n \rangle}\right), \quad (\text{A379})$$

which is arbitrarily large for small mean count. This makes sense: rare events have an arbitrarily large capacity to surprise us and hence to convey information. It is important to note, though, that the maximum entropy per event is a monotonically decreasing function of the mean count. Thus if we are counting spikes from a neuron, counting in larger windows (hence larger mean counts) is always less efficient in terms of bits per spike.

If it is more efficient to count in small time windows, perhaps we should think not about counting but about measuring the arrival times of the discrete events. If we look at a total (large) time interval $0 < t < T$, then we will observe arrival times t_1, t_2, \dots, t_N in this interval; note that the number of events N is also a random variable. We want to find the distribution $P(t_1, t_2, \dots, t_N)$ that maximizes the entropy while holding fixed the average event rate. We can write the entropy of the distribution as a sum of two terms, one from the entropy of the arrival times given the count and one from the entropy of the counting distribution:

$$S = - \sum_{N=0}^{\infty} \int d^N t_n P(t_1, t_2, \dots, t_N) \log_2 P(t_1, t_2, \dots, t_N) \quad (\text{A380})$$

$$= \sum_{N=0}^{\infty} P(N) S_{\text{time}}(N) - \sum_{N=0}^{\infty} P(N) \log_2 P(N), \quad (\text{A381})$$

where we have made use of

$$P(t_1, t_2, \dots, t_N) = P(t_1, t_2, \dots, t_N | N) P(N), \quad (\text{A382})$$

and the (conditional) entropy of the arrival times is given by

$$S_{\text{time}}(N) = - \int d^N t_n P(t_1, t_2, \dots, t_N | N) \log_2 P(t_1, t_2, \dots, t_N | N). \quad (\text{A383})$$

If all we fix is the mean count, $\langle N \rangle = \sum_N P(N) N$, then the conditional distributions for the locations

of the events given the total number of events, $P(t_1, t_2, \dots, t_N | N)$, are unconstrained. We can maxi-

mize the contribution of each of these terms to the entropy [the terms in the first sum of Eq. (A381)] by making the distributions $P(t_1, t_2, \dots, t_N|N)$ uniform, but it is important to be careful about normalization. When we integrate over all the times t_1, t_2, \dots, t_N , we are forgetting that the events are all identical, and hence that permutations of the times describe the same events. Thus the normalization condition is *not*

$$\int_0^T dt_1 \int_0^T dt_2 \cdots \int_0^T dt_N P(t_1, t_2, \dots, t_N|N) = 1, \quad (\text{A384})$$

but rather

$$\frac{1}{N!} \int_0^T dt_1 \int_0^T dt_2 \cdots \int_0^T dt_N P(t_1, t_2, \dots, t_N|N) = 1. \quad (\text{A385})$$

This means that the uniform distribution must be

$$P(t_1, t_2, \dots, t_N|N) = \frac{N!}{T^N}, \quad (\text{A386})$$

and hence that the entropy [substituting into Eq. (A381)] becomes

$$S = - \sum_{N=0}^{\infty} P(N) \left[\log_2 \left(\frac{N!}{T^N} \right) + \log_2 P(N) \right]. \quad (\text{A387})$$

Now to find the maximum entropy we proceed as before. We introduce Lagrange multipliers to constrain the mean count and the normalization of the distribution $P(N)$, which leads to the function

$$\tilde{S} = - \sum_{N=0}^{\infty} P(N) \left[\log_2 \left(\frac{N!}{T^N} \right) + \log_2 P(N) + \lambda_0 + \lambda_1 N \right], \quad (\text{A388})$$

and then we maximize this function by varying $P(N)$. As before the different N s are not coupled, so the optimization conditions are simple:

$$0 = \frac{\partial \tilde{S}}{\partial P(N)} \quad (\text{A389})$$

$$= - \frac{1}{\ln 2} \left[\ln \left(\frac{N!}{T^N} \right) + \ln P(N) + 1 \right] - \lambda_0 - \lambda_1 N, \quad (\text{A390})$$

$$\ln P(N) = - \ln \left(\frac{N!}{T^N} \right) - (\lambda_1 \ln 2)N - (1 + \lambda_0 \ln 2) \quad (\text{A391})$$

Combining terms and simplifying, we have

$$P(N) = \frac{1}{Z} \frac{(\lambda T)^N}{N!}, \quad (\text{A392})$$

$$Z = \sum_{N=0}^{\infty} \frac{(\lambda T)^N}{N!} = \exp(\lambda T). \quad (\text{A393})$$

This is the Poisson distribution.

The Poisson distribution usually is derived (as in our discussion of photon counting) by assuming that the

probability of occurrence of an event in any small time bin of size $\Delta\tau$ is independent of events in any other bin, and then we let $\Delta\tau \rightarrow 0$ to obtain a distribution in the continuum. This is not surprising: we have found that the maximum entropy distribution of events given the mean number of events (or their density $\langle N \rangle / T$) is given by the Poisson distribution, which corresponds to the events being thrown down at random with some probability per unit time (again, $\langle N \rangle / T$) and no interactions among the events. This describes an ‘ideal gas’ of events along a line (time). More generally, the ideal gas is the gas with maximum entropy given its density; interactions among the gas molecules always reduce the entropy if we hold the density fixed.

If we have multiple variables, x_1, x_2, \dots, x_N , then we can go through all of the same analyses as before. In particular, if these are continuous variables and we are told the means and covariances among the variables, then the maximum entropy distribution is again a Gaussian distribution, this time the appropriate multidimensional Gaussian. This example, like the other examples so far, is simple in that we can give not only the form of the distribution but we can find the values of the parameters that will satisfy the constraints. In general this is not so easy: think of the Boltzmann distribution, where we would have to adjust the temperature to obtain a given value of the average energy, but if we can give an explicit relation between the temperature and average energy for any system then we have solved almost all of statistical mechanics!

[obviously this needs to be much better!] One important example is provided by binary strings. If we label 1s by spin up and 0s by spin down, the binary string is equivalent to an Ising chain $\{\sigma_i\}$. Fixing the probability of a 1 is the same as fixing the mean magnetization $\langle \sigma_i \rangle$. If, in addition, we specify the joint probability of two 1s occurring in bins separated by n steps (for all n), this is equivalent to fixing the spin–spin correlation function $\langle \sigma_i \sigma_j \rangle$. For simplicity, consider the case where the system is translation invariant, so the average magnetization is the same at all sites and the correlation function $\langle \sigma_i \sigma_j \rangle$ depends only on $i-j$. The maximum entropy distribution consistent with these constraints is an Ising model,

$$P\{\{\sigma_i\}\} = \frac{1}{Z} \exp \left[-h \sum_i \sigma_i - \sum_{ij} J(i-j) \sigma_i \sigma_j \right]; \quad (\text{A394})$$

note that the interactions are pairwise (because we fix only a two–point function) but not limited to near neighbors. Obviously the problem of finding the exchange interactions which match the correlation function is not so simple.

Another interesting feature of the Ising or binary string problem concerns higher order correlation functions. If we have continuous variables and constrain the two–point

correlation functions, then the maximum entropy distribution is Gaussian and there are no nontrivial higher order correlations. But if the signals we observe are discrete, as in the sequence of spikes from a neuron, then the maximum entropy distribution is an Ising model and this model makes nontrivial predictions about the multipoint correlations. In particular, if we record the spike trains from K separate neurons and measure all of the pairwise correlation functions, then the corresponding Ising model predicts that there will be irreducible correlations among triplets of neurons, and higher order correlations as well [Schneidman et al 2006].

[Where did this come from?] Before closing the discussion of maximum entropy distributions, note that our simple solution to the problem, Eq. (A361), might not work. Taking derivatives and setting them to zero works only if the solution to our problem is in the interior of the domain allowed by the constraints. It is also possible that the solution lies at the boundary of this allowed region. This seems especially likely when we combine different kinds of constraints, such as trying to find the maximum entropy distribution of images consistent both with the two-point correlation function and with the histogram of intensity values at one point. The relevant distribution is a 2D field theory with a (generally nonlocal) quadratic ‘kinetic energy’ and some arbitrary local potential; it is not clear that all combinations of correlations and histograms can be realized, nor that the resulting field theory will be stable under renormalization; the empirical histograms of local quantities in natural images *are* stable under renormalization [Ruderman and Bialek 1994]. There are many open questions here.

[why?]

9. Measuring information transmission

When we study classical mechanics, we can make a direct connection between the positions and momenta that appear in the equations of motion and the positions and momenta of the particles that we “see,” as in the planetary orbits. This connection is a little bit subtle, since we don’t actually measure particle positions; more likely we count the photons arriving at some detector, forming an image, or we measure the delay in propagation of a pulse used in radar, or But one can think of classical mechanics, in contrast to quantum mechanics, as being the domain of physics in which these subtleties are not important. When we move to statistical physics the connection between what we write in equations and what we observe in the world becomes more abstract. The fundamental objects in statistical physics are probability distributions, and as a matter of definition one *cannot* measure a distribution. Instead, Nature (or even a controlled experiment) provides us with samples taken out of these distributions. This has very serious consequences

for any attempt to “measure” information flow.

In thermodynamics, entropy changes are connected to heat flow, and so we can at least measure the difference in entropy between two states by tracking these heat flows. Indeed, there is a long tradition of integrating these changes from some convenient reference to “measure” the entropy of states at intermediate temperatures. As far as I know, there is no analog of this in the information theoretic context. Thus, although Shannon tells us that the entropy is a fundamental property of the distribution out which signals are drawn, there is no universal entropy meter.




FIG. 175 The sampling problem in entropy estimation. At left, the frequency of occurrence found from five examples of $N = 100$ samples drawn out of $K = 10$ bins; the true probability distribution is flat, $p_i = 0.1$ for all i . At right, we estimate the entropy of the distribution by identifying the observed frequencies with probabilities. The distribution of entropies obtained in this way, from many “experiments” with $N = 100$ and $K = 10$, is shown as a solid line, and should be compared with the true entropy, shown by a dashed line at $S_{\text{true}} = \log_2(10)$.

To get a feeling for the problem, consider Fig 175. Here we have a variable that can take on ten possible values ($i = 1, 2, \dots, 10$), all equally likely ($p_i = 0.1$ for all i), and we draw $N = 100$ samples. If we look at the frequency with which each possibility occurs, of course we don’t see an exactly flat distribution. Since with 10 bins and 100 samples we expect 10 samples per bin, it’s not surprising that the fluctuations are on the scale of $1/\sqrt{10} \sim 30\%$. These fluctuations, however, are random—they average to zero if we do the same experiment many times. The problem is that if we identify the frequencies of occurrence as our best estimates of the underlying probabilities, and use these estimates to compute the entropy, we make a systematic error, as is clear from the results in the right panel of Fig 175.

Problem 195: Experiment with sampling. Generate the analog of Fig 175 but with different values for the number of possible values K , where $i = 1, 2, \dots, K$. You should also try different probability distributions (e.g., $p_i \propto 1/i$, Zipf's law). Experiment. Convince yourself that, by identifying probabilities with the observed frequencies of occurrence, you always underestimate the entropy.

The problem illustrated in Fig 175 might seem very specific to the conditions of that simulation (e.g., that the true distribution is flat, and hence the entropy is maximal, so perhaps all errors have to be biased downward?), but in fact is very general. Let's consider drawing samples out of a discrete set of possibilities, $i = 1, 2, \dots, K$, with probabilities $\mathbf{p} \equiv \{p_1, p_2, \dots, p_K\}$. If we draw N samples all together, we will find n_i examples of the outcome i , and of course on average $\langle n_i \rangle = Np_i$. Since we're

counting random events, we expect that the variance of the number of events of type i will be equal to the mean, $\langle (\delta n_i)^2 \rangle = \langle n_i \rangle = Np_i$. If we define the frequency of events in the usual way as $f_i = n_i/N$, then we have

$$\langle f_i \rangle = p_i \quad \text{and} \quad \langle (\delta f_i)^2 \rangle = \frac{p_i}{N}. \quad (\text{A395})$$

But if we identify frequencies as our best estimate of probabilities (and we'll see below in what sense this familiar identification is correct), we can construct a 'naive' estimate of the entropy,

$$S_{\text{naive}} = - \sum_{i=1}^K f_i \log_2 f_i. \quad (\text{A396})$$

Since the frequencies are close to the true probabilities when the number of samples is large, we can do a Taylor expansion around the point $f_i = \langle f_i \rangle = p_i$:

$$\begin{aligned} S_{\text{naive}} &= - \sum_{i=1}^K f_i \log_2 f_i \\ &= - \sum_{i=1}^K (p_i + \delta f_i) \log_2 (p_i + \delta f_i) \end{aligned} \quad (\text{A397})$$

$$= - \sum_{i=1}^K p_i \log_2 p_i - \sum_{i=1}^K \left[\log_2 p_i + \frac{1}{\ln 2} \right] \delta f_i - \frac{1}{2} \sum_{i=1}^K \left[\frac{1}{(\ln 2)p_i} \right] (\delta f_i)^2 + \dots \quad (\text{A398})$$

The first term in the series is the true entropy. The second term is a random error which averages to zero. The third term, however, has a nonzero mean, since it depends on the square of the fluctuations δf_i . Thus when we compute the average of our naive entropy estimate we find

$$\langle S_{\text{naive}} \rangle = S_{\text{true}} - \frac{1}{2 \ln 2} \sum_{i=1}^K \frac{\langle (\delta f_i)^2 \rangle}{p_i} + \dots \quad (\text{A399})$$

$$= S_{\text{true}} - \frac{1}{2 \ln 2} \sum_{i=1}^K \frac{p_i}{N p_i} + \dots \quad (\text{A400})$$

$$= S_{\text{true}} - \frac{K}{2 \ln 2 N} + \dots \quad (\text{A401})$$

Thus, no matter what the underlying true distribution, identifying frequencies with probabilities leads to a *systematic* (not random!) underestimate of the entropy, and the size of this systematic error is proportional to the number of accessible states (K) and inversely proportional to the number of samples (N).

The fact that the systematic errors have a very definite structure suggests that we should be able to correct them. Let us see what happens to our entropy estimates in the

"experiment" of Fig 175 as we change the number of samples N . More precisely, suppose we have only the $N = 100$ samples, but we choose $n < 100$ points out of these samples, and estimate the entropy based only on this more limited data. Equation (A401) suggests that if we plot our entropy estimate vs. $1/n$, we should see a straight line; a higher order version of the same calculation shows that there are quadratic corrections. Indeed, as shown in Fig 176, this works. It is important to note that, for all the accessible range of sample sizes, the entropy estimate is smaller than the true entropy, and this error is larger than our best estimate of the error bar; this really is dangerous. On the other hand, once we recognize the systematic dependence of the entropy estimate on the number of samples, we can extrapolate to recover an estimate that is correct within error bars. What we have seen here about entropy is also true about information, which is a difference between entropies.

It is also important to show that this extrapolation procedure works also for real data, not just for the idealized case where we choose samples independently out of a known distribution. **Decide what examples to use. One from neurons, one from genes, one from sequences?**




FIG. 176 Entropy vs. number of samples. Starting with $N = 100$ samples, as in Fig 175, we draw smaller numbers of samples at random, compute the entropy, and search for the systematic behavior predicted in Eq (A401). Green points are from different subsamplings, blue circles show means and standard deviations. Red line is a linear fit for $n > N/2$, and red dashed line is a quadratic fit to all of the data shown. Red square is the extrapolation, with an error bar $\sqrt{2}$ smaller than the standard deviation found empirically at $n = N/2$, and the dashed black line is $S_{\text{true}} = \log_2(10)$.

So far one figure from neurons, Fig 177.

One might worry that entropy estimates based on extrapolations are a bit heuristic. If we can really convince ourselves that we see a clean linear dependence on $1/N$, things are likely to be fine, but this leaves room for considerable murkiness. Also, since the expansion of the entropy estimate in powers of $1/N$ obviously is not fully convergent, there is always the problem of choosing the regime over which the asymptotic behavior is observed, a widespread problem in fitting to such asymptotic series. While for many purposes these problems can be dismissed, it would be nice to do better. It also is an interesting mathematical challenge to ask if we can estimate the entropy of a probability distribution even when the number of samples we have seen is small, perhaps even smaller than the number of possible states for the system.

Whenever we do a Monte Carlo simulation of a physical system in thermal equilibrium, we are in the “undersampled” limit, where the number of samples we collect must be much smaller than the number of possible states. Usually if we want to estimate entropy from Monte Carlo, we use the identity which relates entropy to an integral of the heat capacity, since heat capacity is related to en-

ergy fluctuations and these are easy to compute at each temperature. Of course, if you just have samples of the state of the system, and don’t actually know the Hamiltonian, you can’t compute the energy and so this doesn’t work. Ma suggested another approach, asking how often the system revisits the same state. In the simple case (relevant for the microcanonical ensemble) where all K possible states are equally likely, the probability that two independent samples are in the same state is $1/K$. But if we have N samples, we have $\sim N^2$ pairs that we can test. Thus we can get a good estimate of the probability of occupying the same state once we observe $N \sim \sqrt{K}$ independent samples, far less than the number of states. As an illustration, Fig 178 shows the frequency of coincidences when we draw N samples from a uniform distribution with $K = 100$ states.

We recall the classic problem of how many people need to be in the room before there is a good chance of two people have the same birthday. The answer is not 365, but more nearly $\sqrt{365}$. Put another way, if we didn’t know the length of the year, we could estimate this by polling people about their birthdays, and keeping track of coincidences. Long before we have sampled all possible birthdays, Fig 178 shows us that our estimate of this coincidence probability will stabilize—which birthdays are represented will vary from sample to sample, but the fraction of coincidences will vary much less.

In these simple examples, the probability distribution is uniform, and so the entropy is just the log of the num-




FIG. 177 Entropy extrapolation with real neural data, from Strong et al (1998a). From the experiment on fly motion-sensitive neurons discussed in Figs 132 and 133, we look at 10-letter words with time resolution $\Delta\tau = 3$ ms. The main figure shows the “Zipf plot” of frequency vs. rank from the full data set. Note that since there are sometimes (but rarely) two spike in one 3 ms bin, there are more than 1024 words. The inset shows the estimated entropy as a function of the (inverse) fraction of the full data set used. The line through the data is from Eq (A401).




FIG. 178 Estimating coincidence probability. Samples are drawn from a distribution that is uniform over $K = 100$ possible states. Green dots show examples, blue circles the mean and standard deviation across many draws of N samples, and black dashed line is the exact answer. We see that the estimate is quite good even when $N \sim \sqrt{K} \ll K$.

ber of possible states, and this in turn is inversely proportional to the probability of a coincidence. So, being able to estimate this probability is equivalent to being able to estimate the entropy. Thus we should be able to generate reliable entropy estimates even in the under-sampled regime, just by counting coincidences. This is a beautiful idea. The challenge is to generalize this idea to non-uniform distributions.

A better understanding of the entropy estimation problem has come through a Bayesian approach. Rather than identifying frequencies with probabilities, we imagine that the distribution itself is drawn from a distribution. To be formal, let the possible states of the system be $i = 1, 2, \dots, K$, and let the probability distribution over these states be $p_1, p_2, \dots, p_K \equiv \mathbf{p}$. This distribution itself is drawn from some distribution function $\mathcal{P}(\mathbf{p})$. The distribution has to be normalized, but it is tempting to think that, other than normalization, all distributions should be equally likely, so that

$$\mathcal{P}(\mathbf{p}) = \frac{1}{Z} \delta \left(\sum_{i=1}^K p_i - 1 \right). \quad (\text{A402})$$

If we observe n_1 samples in the first state, n_2 samples in the second state, and so on, then the probability of this

occurring assuming some distribution \mathbf{p} is

$$P(\{n_i\}|\mathbf{p}) \propto \prod_{i=1}^K p_i^{n_i}, \quad (\text{A403})$$

and so by Bayes' rule we have

$$\mathcal{P}(\mathbf{p}|\{n_i\}) = \frac{P(\{n_i\}|\mathbf{p})\mathcal{P}(\mathbf{p})}{P(\{n_i\})} \quad (\text{A404})$$

$$\propto \frac{1}{Z} \left(\prod_{i=1}^K p_i^{n_i} \right) \delta \left(\sum_{i=1}^K p_i - 1 \right). \quad (\text{A405})$$

If we want to compute our best estimate of the distribution, we have to do the integral

$$\hat{p}_i = \frac{1}{Z} \int d^K \mathbf{p} p_i^{n_i+1} \left(\prod_{j \neq i} p_j^{n_j} \right) \delta \left(\sum_{j=1}^K p_j - 1 \right), \quad (\text{A406})$$

where the normalization Z is given by

$$Z = \int d^K \mathbf{p} \left(\prod_{j=1}^K p_j^{n_j} \right) \delta \left(\sum_{j=1}^K p_j - 1 \right). \quad (\text{A407})$$

To make progress we introduce the Fourier representation of the delta function, so that, for example,

$$Z = \int d^K \mathbf{p} \left(\prod_{j=1}^K p_j^{n_j} \right) \int \frac{d\lambda}{2\pi} \exp \left(+i\lambda \sum_{j=1}^K p_j - i\lambda \right) \quad (\text{A408})$$

$$= \int \frac{d\lambda}{2\pi} e^{-i\lambda} \prod_{j=1}^K \int dp_j p_j^{n_j} e^{i\lambda p_j}. \quad (\text{A409})$$

Since we have the delta function, we are free to let the integrals over p_j run from 0 to ∞ ; the delta function will enforce the constraint that $p_i \leq 1$ for all i . Then the key ingredient of the calculation, then, is the integral

$$f(n; \lambda) = \int_0^\infty dp p^n e^{i\lambda p}. \quad (\text{A410})$$

At the end of our calculation we will have to integrate over λ . Let's assume that we will be able to deform the contour of this integral into the complex λ plane in such a way that the p integral in Eq (A410) is well behaved. Then we can write

$$f(n; \lambda) = \int_0^\infty dp p^n e^{i\lambda p} \quad (\text{A411})$$

$$= \int_0^\infty dp p^n e^{-(i\lambda)p} = \frac{n!}{(-i\lambda)^{n+1}}. \quad (\text{A412})$$

Putting these pieces together, we have

$$\mathcal{Z} = \int \frac{d\lambda}{2\pi} e^{-i\lambda} \prod_{j=1}^K \frac{n_j!}{(-i\lambda)^{n_j+1}} \quad (\text{A413})$$

$$= \left(\prod_{j=1}^K n_j! \right) \int \frac{d\lambda}{2\pi} \frac{e^{-i\lambda}}{(-i\lambda)^{\sum_{j=1}^K (n_j+1)}} \quad (\text{A414})$$

$$= \left(\prod_{j=1}^K n_j! \right) \int \frac{d\lambda}{2\pi} \frac{e^{-i\lambda}}{(-i\lambda)^{N+K}}, \quad (\text{A415})$$

where $N = \sum_j n_j$ is the total number of samples, and as before K is the number of possible states. A similar argument gives

$$\hat{p}_i = \frac{1}{\mathcal{Z}} (n_i + 1)! \left(\prod_{j \neq i} n_j! \right) \int \frac{d\lambda}{2\pi} \frac{e^{-i\lambda}}{(-i\lambda)^{N+K+1}} \quad (\text{A416})$$

$$= \frac{(n_i + 1)! \left(\prod_{j \neq i} n_j! \right)}{\prod_{j=1}^K n_j!} \times \frac{\int \frac{d\lambda}{2\pi} \frac{e^{-i\lambda}}{(-i\lambda)^{N+K+1}}}{\int \frac{d\lambda}{2\pi} \frac{e^{-i\lambda}}{(-i\lambda)^{N+K}}} \quad (\text{A417})$$

$$= (n_i + 1) \frac{\int \frac{d\lambda}{2\pi} \frac{e^{-i\lambda}}{(-i\lambda)^{N+K+1}}}{\int \frac{d\lambda}{2\pi} \frac{e^{-i\lambda}}{(-i\lambda)^{N+K}}}. \quad (\text{A418})$$

Thus, $\hat{p}_i \propto n_i + 1$, so to get the normalization right we must have

$$\hat{p}_i = \frac{n_i + 1}{N + K}. \quad (\text{A419})$$

This should be contrasted with the naive estimate of probabilities based on counting frequencies, $\hat{p}_i = n_i/N$. The Bayesian estimate, with a ‘flat’ prior on the space of distributions, is equivalent to the naive approach but with one extra count in every bin. This estimate never predicts probability zero, even in states never observed to occur, and is in some sense ‘smoother’ than the frequencies. The trick of adding such pseudocounts to the data goes back, it seems, to Laplace, although I don’t think he had the full Bayesian justification.

Problem 196: Normalization. Derive Eq (A419) directly by doing the integrals in Eq (A418).

What does this have to do with entropy estimation? Somewhat heuristically, it has been suggested that by using different numbers of pseudocounts one can improve the quality of entropy estimation. More deeply, I think, the Bayesian estimate gives us a very different view of *why* we make systematic errors when we try to compute entropies from data. Recall that when we use the naive

identification of frequencies with probabilities, we underestimate the entropy, as in Eq (A401). It is tempting to think that we are underestimating the entropy simply because, in a finite sample, we have not seen all the possibilities. With the Bayesian approach and a flat prior, however, the probability distributions that we estimate are smoother than the true distribution, and correspondingly we expect that the entropy will be overestimated. In fact this is true, but the problem really is more serious than this.

Suppose that we don’t yet have any data. Then all we know is that the probability distribution \mathbf{p} will be chosen out of the distribution $\mathcal{P}(\mathbf{p})$. This seems innocuous, since this distribution is flat and hence presumably unbiased. But we can calculate the average entropy in this distribution,

$$\langle S \rangle_{\text{prior}} \equiv \int d^K \mathbf{p} \left(- \sum_{i=1}^K p_i \log_2 p_i \right) \mathcal{P}(\mathbf{p}), \quad (\text{A420})$$

using the same tricks that we used above, and we find

$$\langle S \rangle_{\text{prior}} = \psi_0(K + 1) - \psi_0(1), \quad (\text{A421})$$

where $\psi_0(x)$ is a polygamma function,

$$\psi_m(x) = \left(\frac{d}{dx} \right)^{m+1} \Gamma(x). \quad (\text{A422})$$

The details of the special functions are not so important. What is important is that, when the number of states K is large,

$$\langle S \rangle_{\text{prior}} = \log_2 K - \mathcal{O}(1). \quad (\text{A423})$$

Thus, although we are choosing distributions from a flat prior, the entropies of these distributions are biased toward the maximum possible value. This bias is actually very strong. The entropy is the average of many terms, and although these terms can’t be completely independent (the probabilities must sum to one), one might expect the central limit theorem to apply here, in which case the fluctuations in the entropy will be $\sigma_S \sim 1/\sqrt{K}$, which for large K is very small indeed. What this means is that the distributions chosen out of $\mathcal{P}(\mathbf{p})$ are overwhelmingly biased toward having nearly maximal entropy. While the prior on the distributions is flat, the prior on entropies is narrowly concentrated around an average entropy which, for large K , is almost $\log_2 K$.

Problem 197: Entropies in a flat prior. Derive the mean and standard deviation of the entropy in the flat prior, $\mathcal{P}(\mathbf{p})$ from Eq (A402). Verify Eq (A423).

Just to make the problem clear, suppose that our system has only two states, as with heads and tails for a coin. Let the probability of heads be q , so that the entropy is

$$S(q) = -q \log_2(q) - (1 - q) \log_2(1 - q). \quad (\text{A424})$$

If we assume that q is chosen from some distribution $\mathcal{P}(q)$, then the distribution of entropies can be found from

$$P(S)dS = \mathcal{P}(q)dq \quad (\text{A425})$$

Since $dS/dq = 0$ at the point where $S = 1$ bit, the distribution $P(S)$ must be singular there unless the prior on q itself has a compensating singularity. Thus, a prior which is flat in q is strongly biased in S . The situation is even worse for systems with many states, because of phase space considerations: if we want to have low a low entropy distribution, then many of the p_i must be confined to very small values, and this means that the volume in \mathbf{p} space associated with low entropy is small. While only one distribution has precisely the maximum entropy, there are many distributions that are close.

Problem 198: A flat prior on S . Show that, for the problem of coin flips, having a flat prior on the entropy S is equivalent to a prior

$$\mathcal{P}(q) = \left| \log_2 \left(\frac{q}{1-q} \right) \right|. \quad (\text{A426})$$

If we flip a coin N times and observe n heads, then Bayes' rule tell us that

$$\mathcal{P}_N(q|n) \propto \mathcal{P}(q)q^n(1-q)^{N-n}, \quad (\text{A427})$$

and we can use this to estimate the entropy

$$\hat{S}(n, N) = \int_0^1 dq \mathcal{P}_N(q|n) [-q \log_2(q) - (1 - q) \log_2(1 - q)]. \quad (\text{A428})$$

(a.) For $N = 10$, plot $\hat{S}(n, N)$ vs. n . Compare your results with the naive estimate,

$$S_{\text{naive}}(n, N) = -\frac{n}{N} \log_2 \left(\frac{n}{N} \right) - \left(1 - \frac{n}{N} \right) \log_2 \left(1 - \frac{n}{N} \right). \quad (\text{A429})$$

(b.) Suppose that you are actually flipping a coin in which the probability of heads is $q_{\text{true}} \neq 1/2$. Simulate N such flips, and use your results to estimate the entropy according to both Eq's (A428) and (A429). How do these estimators evolve as a function of N ? Hints: Remember that we have seen the results for the naive estimator already, and that since this is a small system (only two states) the interesting behavior is at smaller N .

[Add figure on entropy estimation for binary variables, with flat priors on entropy or probability.] All of this suggests that we could do a much better job of entropy estimation in a Bayesian framework where the $\mathcal{P}(\mathbf{p})$ is chosen to be flat is S . I don't know of anyone who has given a complete solution to this problem. A partial solution has been proposed by noticing that there is a well

known generalization of the flat prior, the Dirichlet family of priors

$$\mathcal{P}_\beta(\mathbf{p}) = \frac{1}{Z(\beta)} \left(\prod_{i=1}^K p_i^{\beta-1} \right) \delta \left(\sum_{i=1}^K p_i - 1 \right). \quad (\text{A430})$$

Evidently the flat prior corresponds to $\beta = 1$, and this is biased toward large entropies, as we have seen. As β gets smaller, the average entropy $\bar{S}(\beta)$ of a distribution drawn out of $\mathcal{P}_\beta(\mathbf{p})$ gets smaller, but for each value of β the distribution of entropies remains quite narrow. This suggests that if we form the prior

$$\mathcal{P}(\mathbf{p}) = \int_0^1 d\beta \left| \frac{d\bar{S}(\beta)}{d\beta} \right|^{-1} \mathcal{P}_\beta(\mathbf{p}), \quad (\text{A431})$$

it will be approximately flat in entropy. This seems to work, although it is computationally intensive. As far as I know it gives the best results of any estimation procedure so far in, for example, the analysis of neural spike trains. If we dig into the integrals that define the entropy estimate, it turns out that the key pieces of data are coincidences, in which more than one sample falls into the same bin, and in this sense we seem to have found a generalization of Ma's ideas to non-uniform distributions.

Problem 199: One more problem about entropy estimation. [make up one more?]

Again it is important to ask whether these ideas actually work with real data. In experiments on the motion-sensitive neuron H1 in the fly visual system, we can in many cases collect enough data to sample the underlying distributions of neural responses, so we have ground truth. At the same time, we can look only at a small fraction of these data and ask how well our estimation procedure works. An example is shown in Fig 179.

Is there more to say? Or need more details in things already said?

The idea that naive counting leads to systematic errors in entropy estimation goes back, at least, to Miller (1955). The importance of this for the analysis of information transmission in neurons was emphasized by Treves and Panzeri (1995), who also brought more sophistication to the calculation of the series expansion that we have started here. Shortly after this, Strong et al (1998a) showed how these extrapolation methods could be used to estimate entropy and information in neural responses to complex, dynamic sensory inputs. An important technical point is that Strong et al took seriously the $1/N$ behavior of the entropy estimate, but didn't use an analytic calculation of the slope of S_{est} vs $1/N$; the reason is




FIG. 179 Estimating entropies at one slice of time in the neural response to naturalistic stimuli, from Nemenman et al (2004). Neural responses are discretized with $\Delta\tau = 2$ ms resolution, and we look at 8-letter words. The stimulus is motion outdoors, and the motion is repeated many times; here we focus on the distribution of responses at one moment relative to this repeat, for which we can collect up to 196 samples from the repetitions. The open symbols show the “naive” or maximum likelihood estimate in which we identify the observed frequencies with probabilities and plug in to the computation of entropy. As expected, this estimate has a significant dependence on the number of samples, but extrapolates smoothly according to Eq (A401). In contrast, the NSB estimator based on the prior in Eq (A431) remains constant within error bars, always agreeing with the extrapolation.

that some seemingly possible neural responses are expected to have probability zero, because there is a hard core repulsion (“refractoriness”) between spikes, but we don’t know in advance exactly how big this effect will be. As a result, the actual number of possible states K is uncertain, and in addition it is not true that all the samples collected in the experiment will be independent. Both these effects leave the $1/N$ behavior intact, but change the slope.

Miller 1955: Note on the bias of information estimates. GA Miller, in *Information Theory in Psychology: Problems and Methods II–B*, H Quastler, ed pp. 95–100 (Free Press, Glencoe IL, 1955).

Strong et al 1998a: Entropy and information in neural spike trains. SP Strong, R Koberle, RR de Ruyter van Steveninck & W Bialek, *Phys Rev Lett* **80**, 197–200 (1998).

Treves & Panzeri 1995: The upward bias in measures of information derived from limited data samples. A Treves & S Panzeri, *Neural Comp* **7**, 399–407 (1995).

The idea of estimation entropy by counting coincidences is presented in Ma (1981), although it has precursors in Serber (1973). The attempts to build a ‘flat prior’ on the entropy are described in a series of papers by Nemenman (2002) and co-workers (Nemenman et al 2002). Moments of the entropy distribution in Dirichlet priors were calculated by Wolpert & Wolf (1995), so you can see where the polygamma functions come from. Within the family of Dirichlet priors, we have noted that the flat prior on distributions ($\beta = 1$) was discussed by Laplace (1814) as the idea of starting with one pseudocount in every bin; Jeffreys (1946), and later Krichevskii & Trofimov (1981) proposed using half a pseudocount ($\beta = 1/2$),

while Schurmann & Grassberger (1996) suggested that entropy estimates could be improved by adjusting the number of pseudocounts to the number of bins, $\beta = 1/K$. The procedure developed by Nemenman et al integrates over all β , allowing a dominant β to emerge that is matched to the structure of the data and to the number of samples, as well as to K . The example in Fig 179 is from Nemenman et al (2004).

Jeffreys 1946: An invariant form for the prior probability in estimation problems. H Jeffreys, *Proc R Soc Lond Ser A* **186**, 453–461 (1946).

Krichevskii & Trofimov 1981: The performance of universal encoding. R Krichevskii & V Trofimov, *IEEE Trans Inf Thy* **27**, 199–207 (1981)

Laplace 1814: *Essai philosophique sur les probabilités* (Courcier, Paris, 1814). *A Philosophical Essay of Probabilities*, translated by F Truscott & F Emory (Dover, New York, 1951).

Ma 1981: Calculation of entropy from data of motion. S–K Ma, *J Stat Phys* **26**, 221–240 (1981).

Nemenman 2002: Inference of entropies of discrete random variables with unknown cardinalities. I Nemenman, arXiv:physics/0207009 (2002).

Nemenman et al 2002: Entropy and inference, revisited. I Nemenman, F Shafee & W Bialek, in *Advances in Neural Information Processing 14*, TG Dietterich, S Becker & Z Ghahramani, eds, pp 471–478 (MIT Press, Cambridge, 2002); arXiv:physics/0108025 (2001).

Nemenman et al 2004: Entropy and information in neural spike trains: Progress on the sampling problem. I Nemenman, W Bialek & R de Ruyter van Steveninck, *Phys Rev E* **69**, 056111 (2004); arXiv:physics/0306063 (2003).

Serber 1973: *Estimation of Animal Abundance and Related Parameters*. GAF Serber (Griffin, London, 1973).

Schurmann & Grassberger 1996: Entropy estimation of symbol sequences. T Schurmann & P Grassberger, *Chaos* **6**, 414–427 (1996).

Victor 2002: Binless strategies for estimation of information from neural data. J Victor *Phys Rev E* **66**, 051903 (2002).

Wolpert & Wolf 1995: Estimating functions of probability distributions from a finite set of samples. DH Wolpert & DR Wolf, *Phys Rev E* **52**, 6841–6854 (1995).

The specific problem of estimating entropy in neural responses has the added feature that spikes are discrete, but they can occur at any time, so there is a question of whether we should view the whole problem as discrete (with bins along the time axis) or continuous (with a metric along the time axis); the discussion here has focused on discrete problems. For metric space approaches to spike trains, see Victor (2002). For a general overview of the entropy estimation problem, with particular attention to the challenges posed by neural data, see Paninski (2003), who tries to make many of the heuristic arguments in the field more rigorous. One can also view entropy estimation as a problem in computational complexity—how many samples, and hence how many computational steps, do we need in order to approximate the entropy to some level of accuracy? For an approach in this spirit, see Batu et al (2002). [\[One more recent paper comparing the effectiveness of different approaches\]](#)

Batu et al 2002: The complexity of approximating the entropy. T Batu, S Dasgupta, R Kumar & R Rubinfeld, in *Proc 34th Symp Theory of Computing (STOC)*, pp 678–687 (2002).

Paninski 2003: Estimation of entropy and mutual information. L Paninski, *Neural Comp* **15**, 1191–1253 (2003).

Victor 2002: Binless strategies for estimation of information from neural data. J Victor *Phys Rev E* **66**, 051903 (2002).