

February 11, 2004

In order to survive in the world, organisms have to solve a wide variety of problems, and in many cases these are physics problems: capturing energy in one form and converting it to another, sensing and processing small signals in the environment, building reliable machines on a microscopic scale in the face of thermal fluctuations, More subtly, the successful organism must act in ways that make sense in relation to events in the outside world—behavior must be “functional”—and this means that the mechanisms of its behavior must at least implicitly embody a model of the surrounding physical world. The physics problems that arise in the life of an organism thus range from the dynamics of single molecules to learning about the dynamics of the environment, and I have been interested in a collection of topics that span this range of issues; the list here includes some favorite papers. The interface between physics and biology provides a fantastic opportunity for interaction between theory and experiment, and so half of the papers here were written jointly with experimentalists.

Note: Numbers refer to a full publication list, available at http://www.princeton.edu/~wbialek/publications_wbialek.html

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

From a physicist’s point of view, some of the most remarkable facts about biological systems concern the absolute sensitivity of the the sense organs: eyes that count single photons, ears that detect vibrations comparable in amplitude to Brownian motion, chemical sensors that count single molecules, thermal sensors with milliKelvin sensitivity at room temperature, In this work I tried to bring these very different observations together, combining classic data with some of my own arguments about more recent experiments. The goal was to see performance near the physical limits as a general principle. I also wanted to address the issue of what happens after the detectors: while one might concede that physics is relevant for the design of receptor cells, surely once the resulting signals enter the brain one is doing biology (or even psychology) and not physics, or so go the prevailing prejudices. Against these prejudices I offered several examples where the nervous system does nontrivial computations and still reaches answers whose precision is set by physical principles. This broader notion of optimal performance, pushing the physical limits wherever possible, remains a radical view of biology in general and of neural computation in particular.

[21.] 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).

All of the sensory signals reaching the brain are encoded in sequences of identical, discrete pulses called action potentials or spikes. Spike trains in other regions of the brain can represent motor commands or, more subtly, plans and intentions, and there is no reason to doubt that even our private thoughts are represented in this way. The way in which spikes represent the external world has been studied since the 1920s. In this paper we argued that one should take a new point of view on this problem. Instead of asking how known signals in the outside world are encoded in the average properties of spike trains, we asked how the brain—which has only the spike trains to work with—could make inferences about unknown sensory stimuli. Specifically, we showed how to characterize the distribution of sensory inputs that are consistent with the neural response, thus quantifying the (un)certainty of inferences from the spike train. We introduced information theoretic ideas to the analysis of spike trains, and showed, for example, how one could measure the information conveyed by the absence of spikes. These ideas were used in the design and analysis of experiments on the fly visual system, specifically a neuron H1 that is responsible for extracting information about (horizontal) rigid body motion across the whole visual field. There are several specific points that have become important in recent work: the demonstration that short sequences of spikes are informative only about projections of the stimulus onto spaces of low dimensionality, that similar spike trains stand for similar stimulus waveforms, and that patterns of spikes can convey more information than expected by summing the contributions of individual spikes. Beyond these specific results, I believe that the point of view expressed in this paper set the agenda for much of what followed.

[25.] Coding and computation with neural spike trains. W Bialek & A Zee, *J. Stat. Phys.* **59**, 103–115 (1990).

Inspired in part by the results in the fly, we set out to study the problem of coding and decoding in simple models of spiking neurons. Probably the most important result was that there is a large regime in which signals can be *decoded* by linear (perturbative) methods even though the encoding is strongly nonlinear. The small parameter that makes this work is the mean number of spikes per correlation time of the signal, suggesting that spike trains can be decoded linearly if they are sparse. Some further aspects of the theory were developed in [29]. But even the simplest calculations suggested an interesting experiment: could we decode the spike trains of real neurons?

[34.] Reading a neural code. W Bialek, F Rieke, RR de Ruyter van Steveninck, & D Warland, *Science* **252**, 1854–1857 (1991).

Returning to the experiments, we showed that it really is possible to decode neural spike trains and thereby reconstruct the time dependent sensory stimulus. The coexistence of linear decoding with nonlinear encoding was implicit in this work, but made explicit in further analysis described in [65]. This work was intended as a proof of principle, but

we also found that the reconstructions were surprisingly precise: errors corresponded to $\sim 0.06^\circ$ over ~ 30 ms, which is $\sim 20\times$ smaller than the lattice spacing of detectors in the compound eye or $\sim 10\times$ smaller than the nominal ‘diffraction limit’ due to blur by the eye’s lenses. Resolution beyond the sampling and diffraction scales is also known in human vision, and the collection of perceptual phenomena in which this occurs is called hyperacuity.¹ This led us to wonder about the physical limits to motion estimation—blur due to diffraction through the lenses of the compound eye and noise due to the random arrival of photons at the receptors. In fact the observed performance is very close to this limit, so that even four layers of neurons away from the receptors it is still the physics of the inputs that sets the precision of computation.² The ideas of decoding or stimulus reconstruction caught on among experimentalists, and have since been applied to systems ranging from motor control in crabs to visual motion perception in monkeys; even some current techniques for decoding information in primate motor cortex—with the hope of driving prosthetic devices—are straightforward outgrowths of the reconstruction methods first used in this paper.

[37.] Vibrationally enhanced tunneling as a mechanism for enzymatic hydrogen transfer. WJ Bruno & W Bialek, *Biophys. J.* **63**, 689–699 (1992).

The transfer of hydrogen atoms or ions is central to a wide variety of biological processes. There has long been interest in the possibility that these reaction proceed by quantum tunneling, but the evidence was murky at best. In many enzymes the hydrogen transfer reactions show anomalous isotope effects, as expected for tunneling, but also substantial temperature dependencies, as expected for thermal activation. In this work we considered a scenario in which (classical) protein motions could enhance the quantum tunneling of hydrogen in enzymatic reactions by causing fluctuations in the shape of the barrier. In the semiclassical limit we showed that almost any reasonable model in this class leads to a very simple phenomenology that fits the puzzling pattern of temperature dependent kinetic isotope effects. New experiments have provided yet stronger evidence for the importance of protein enhanced tunneling, and current theoretical activity is concentrated on detailed quantum chemical calculations that embody the basic scenario proposed in our work. I found it very pleasing to see how biology could exploit this interplay between classical and quantum dynamics.

[38.] 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).

Photosynthesis begins with a photon induced transfer of an electron from one large organic molecule to another, both held in a protein framework. This step is complete in ~ 3

¹For a review of hyperacuity see Section 4.2 of [65]. Perceptual hyperacuity usually is demonstrated in tasks that involve discrimination among discrete alternative signals; the reconstruction experiments allowed the demonstration of comparable precision in a more natural task of continuous estimation. Experiments that are more analogous to the discrimination experiments have also been done on the H1 neuron [57], and a preliminary account of these experiments (in 1984) may have been the first report of hyperacuity in the responses of a single neuron.

²For details of the limits to motion estimation, see [29, 42].

psec, one of the fastest reactions of its type ever observed. Chemistry as we usually think about it operates in the limit where reactions are slow compared with internal molecular relaxation rates, and on these grounds alone it seemed unlikely that the initial event of photosynthesis could be thought of as a conventional chemical kinetic process. In more detail, if one tries to estimate the matrix elements among the relevant electronic states and use the golden rule to calculate the transfer rate, there are considerable uncertainties but it seemed hard to get the right answer. In this work we showed that there is a surprisingly broad regime in which electronic matrix elements, vibrational level spacings and relaxation rates are all comparable, so that one can be poised in between the golden rule regime and coherent oscillation. We argued that this regime would be consistent with various aspects of the phenomenology in photosynthesis, and that in a quantitative sense this regime corresponds to the maximum possible transfer rate given the electronic matrix elements. As we were completing the paper, a new set of ultrafast spectroscopy experiments at lower temperatures revealed the coherent oscillations that would occur in our scenario if relaxation rates were reduced. In a similar spirit we studied a simple model for the initial events in the visual pigments [16], combining intuition from condensed matter physics work on conjugated polymers with a (then) novel simulation technique that combined molecular dynamics with diagonalization of a model Hamiltonian for the electrons;³ again subsequent experiments detected the coherence effects expected from the theory, although it would be an overstatement to say that the theory was confirmed. These are some of the only cases I know where quantum coherence really is important in a biological process.

[45.] Coding efficiency and information rates in sensory neurons. F Rieke, D Warland, & W Bialek, *Europhys. Lett.* **22**, 151–156, (1993).

The stimulus reconstruction techniques [25,34] make it possible to estimate the rate at which neurons convey information about the sensory stimulus.⁴ Using these methods, it was shown that single sensory neurons can transmit enormous amounts of information. The record thus far is 300 bits/s in a cricket mechanical sensor. In this system, and in the vibration sensors of the frog, the information rate is within a factor of two of the fundamental physical limit set by the entropy of the spike train itself. This limit was calculated by MacKay and McCulloch in 1952 in the first application of information theory to the neural code; this paper provided the first evidence that the physical limits are relevant to real neurons. In the frog auditory system one can get even closer to the physical limit by using sensory inputs that capture some of the statistical structure of the frog's natural acoustic environment [58].

[53.] Statistical mechanics and visual signal processing. M Potters & W Bialek, *J. Phys. I France* **4**, 1755–1775 (1994).

³In retrospect this is a sort of poor man's Car-Parinello method (and done at the same time), using tight binding rather than density functionals. I think it is quite a powerful technique, and we should have made more of it at the time.

⁴More precisely, the reconstruction method provides a rigorous lower bound on the rate of information transmission [48]; see also Section 3.2.3 of [65].

Inspired by the observation of near optimal performance in the fly’s motion estimation system, we set out to understand the algorithmic requirements for optimal estimation. Conventional approaches involve searching a set of possible strategies for the best within the set, but we showed how one could map the problem of estimation in the presence of noise onto a statistical mechanics problem in which the data act as external fields and the estimator is the expectation value of some order parameter. Estimation theory then is reduced to the computation of (perhaps strongly nonlinear) response functions, and standard approximations in statistical mechanics map to different regimes of the signal processing problem. Applying the general framework to the problem of motion estimation in the fly, we showed that the optimal estimation strategy has very different behaviors in different sensory environments. In particular, the optimal estimator interpolates between popular models for motion estimation, which arise as limiting cases of the full theory. An inevitable prediction of the theory is that the optimal processor must change its strategy, or *adapt* to changes in the statistics of the input signals. Preliminary experiments gave clear evidence of the “statistical adaptation” predicted by the theory [54, 59, 61], and recent experiments provide a direct confirmation of the combination of nonlinear operations predicted by the theory [102]. In the case of motion estimation there is nothing deep about the statistical mechanics problems that we have to solve, but in cases where stimuli have ambiguous interpretations (as in the Necker cube) we found that the estimation problem maps to a random field model and that the nontrivial statistical mechanics of the random field problem really does correlate with the phenomenology of multistable percepts [56]. This seems to me a very clear example of how the idea of optimal performance can generate striking and even counterintuitive predictions, in this case predicting fluctuations in perception even when inputs are constant and the signal to noise ratios are high.

[66.] 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).

There have been fifty years of debate on the question of whether the detailed timing of spikes is important. With precise timing the system would have access to a much greater representational capacity: the entropy of neural responses is higher at higher time resolution, but is this capacity used efficiently? Here we showed for the first time how to measure—without reference to any models or assumptions—the information content of neural spike trains as they encode naturalistic, dynamic inputs. The result was that the motion sensitive neurons in the fly visual system use roughly half of the spike train entropy to carry visual information, and this efficiency is approximately constant from a time resolution of ~ 1 sec down to ~ 1 msec. As in the frog ear, more informative and efficient responses are observed under more natural conditions [63, 77]. This work has opened the door to a more quantitative characterization of the neural code in a wide variety of systems. While in this original work we tried to avoid sophisticated statistical methods, it is clear that further progress on many systems will require inference of probability distributions and entropies in a regime far from complete sampling, a problem of considerable theoretical interest on its own; we have made some progress in this direction [83, 101].

[69.] 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.

When Shannon developed information theory he left open the problem of assigning relevance to a signal. Here we showed that if we observe one signal but are interested in another, then the statistical associations between these signals defines what is relevant and one can (selectively) compress the observed signal to “squeeze out” the relevant bits. Crucially, this formulation does not require any assumptions about what it means for signals to be similar; indeed the various signals need not even live in a metric space. There are deep connections to clustering—especially to the statistical mechanics formulation in which separate clusters emerge through a series of phase transitions—and many different problems from signal processing to learning can be cast into this unified information theoretic framework (a more detailed account is in preparation). We believe that this is a fundamentally new and principled approach to a wide variety of problems of interest both as models of the problems solved by the brain and as practical problems in their own right.

[71.] Synergy in a neural code. N Brenner, SP Strong, R Koberle, W Bialek, & RR de Ruyter van Steveninck, *Neural Comp.* **12**, 1531–1552 (2000).

Timing of spikes could be significant because each spike points precisely to an event in outside world, or because the system really uses temporal patterns of spikes to convey something special. Here we showed how to give this question a model independent information theoretic formulation: do patterns of spikes carry more information than expected by summing the contributions of individual spikes? To answer this requires measuring the information carried by particular candidate symbols in the code, and we show how this can be done with real data, making connections between the information theoretic quantities and the more familiar correlation functions of the spike train; although we focused on patterns across time in a single cell, everything generalizes to patterns across a population of cells. For the fly’s motion sensitive neuron, we do observe synergistic coding, and this synergy is a significant component of the high coding efficiency seen in [66]. Again this work provides the tools to settle an old question, and in one test case provides the positive and interesting answer.

[72.] Adaptive rescaling optimizes information transmission. N Brenner, W Bialek, & R de Ruyter van Steveninck, *Neuron* **26**, 695–702 (2000).

The direct demonstration of high coding efficiency [45, 66, 77] supports strongly the forty year old idea that the construction of an efficient representation could be the goal of neural computation. Efficient representations must be matched to the statistical structure of the input signals, and it therefore is encouraging that we observe higher coding efficiency for more naturalistic signal ensembles [58, 63, 74], but it usually was assumed that such matching could occur only over the long time scales of development or evolution. In [53, 55] we proposed that adaptation to statistics would occur in real time, to exploit the intermittent structure of natural signals (e.g. [52]), and in [54, 61, 62] we presented evidence

that this occurs both in the fly and in the vertebrate retina. Here we analyzed an example in detail, and found that adaptation to changes in the variance of the input has a striking form, rescaling the input/output relation of the neuron so that signals are coded in relative units. Further, the precise choice of the scaling factor serves to optimize information transmission; this is the first direct demonstration that an optimization principle is at work in the brain. Subsequent work [78] showed that the dynamics of this adaptation itself is optimal, so that the speed with which the system adjusts to a change in input distribution is close to the limit set by the need to gather statistics.

[74.] Stability and noise in biochemical switches. W Bialek, in *Advances in Neural Information Processing 13*, TK Leen, TG Dietterich, & V Tresp, eds., pp. 103–109 (MIT Press, Cambridge, 2001).

Many processes in biology, from the regulation of gene expression in bacteria to the storage of memories in the brain, involve switches constructed from networks of biochemical reactions. Crucial molecules are present in small numbers, raising questions about noise and stability—what are the limits of miniaturization as biology tries to build computational elements from just a few molecules? Here I showed that, perhaps surprisingly, switches stable for years and switchable in milliseconds can be built from just tens of molecules. This has serious implications for how we think about many cellular processes, and should be testable directly, although clearly this is just a small part of a much larger problem. I also used this paper to point out errors in the persistent arguments for the inadequacy of theory (and hence for the necessity of simulation) in the study of these biochemical networks. In [99] we return to the problem of noise in biochemical reactions, trying to define the limits to the precision of signaling and arguing that recent experiments on intracellular signaling and on the regulation of gene expression demonstrate performance close the physical limits which we derive.

[75.] Universality and individuality in a neural code. E Schneidman, N Brenner, N Tishby, RR de Ruyter van Steveninck, & W Bialek, in *Advances in Neural Information Processing 13*, TK Leen, TG Dietterich, & V Tresp, eds., pp. 159–165 (MIT Press, Cambridge, 2001).

One of the major challenges in thinking quantitatively about biological systems is the variability among individuals. In the context of the neural code, we can ask if different animals share similar neural representations of the same sensory inputs. The problem of comparing neural representations is similar in several ways to the problem of comparing DNA sequences, and we argue that rather than using conventional metric or string matching methods one should take a model independent information theoretic approach; we believe that this is a substantial conceptual advance that should have implications back to the bioinformatics problem. We then find that the fly’s visual system has a quantifiable mixture of universality and individuality: what is universal is the efficiency of the code, and what is individual is the precise way in which patterns of spikes are used to achieve this efficiency. In [87] we build on these ideas to classify the population of neurons in the optic nerve.

[76.] Predictability, complexity and learning. W Bialek, I Nemenman, & N Tishby, *Neural*

Comp. **13**, 2409–2463 (2001).

We have reached an understanding of the connections between learning and complexity as unified by the idea of *predictive information*, which is equivalent to subextensive components in the entropy. The results provide a conclusive answer to the long standing problem of how to characterize the complexity of time series, and serve to unify ideas from different areas of physics and computer science. In particular we can classify data streams by their complexity, and if there is something to be learned from the data stream then this classification corresponds to measures for the complexity of the model that can be learned. From a technical point of view it was essential to have a calculable example in the regime where models to be learned cannot be described by a finite number of parameters, and in related work we showed how these nonparametric learning problems could be given a field theoretic formulation [60, 80]. Perhaps the most interesting direction to grow out of this work is the possibility of measuring directly the complexity of models used by humans and other animals as they learn about the world. Together with the concept of selective compression [69], these ideas and methods lead to a whole program of research on the relevant features of natural signals, from DNA sequences and protein structures to communication sounds to language itself [84].

[88.] Computation in single neurons: Hodgkin and Huxley revisited. B Agüera y Arcas, AL Fairhall, & W Bialek, *Neural Comp.* **15**, 1715–1749 (2003).

The combination of techniques from molecular biology, low noise electrical recordings and optical imaging means that we can now make essentially exact models for the dynamics of ion channels and for the electrical dynamics of neurons. These models in fact are quite complex, but they are not understandable as models of *computation*; the situation is similar to having accurate device physics models without knowing about the elementary instructions that the devices should implement. In this paper we formulate a description of the computation performed by a neuron as a combination of dimensional reduction and nonlinearity, and show how one can test this view in simulations of the simplest model neuron. Already this simplest model has dynamics much richer than the simple threshold device of conventional neural networks. The methods we use to analyze models can also be used to analyze experiments on real neurons, and indeed our approach grew out of the dimensionality reduction results in [21] and were crucial in the analysis of adaptation and optimization [72]. I believe that the idea of dimensionality reduction will be much more widely applicable to the description of neural computation [89, 93, 102], and that the question we are addressing here is much deeper than the biological context: given a dynamical system, what does it compute?