Reliability and information transmission in spiking neurons

William Bialek and Fred Rieke

William Bialek and Fred Rieke are at the NEC Research Institute, 4 Independence Way, Princeton, NJ 08540, USA.

Spiking neurons encode continuous, time-varying signals in sequences of identical action potentials. Relatively simple algorithms allow one to 'decode' this neural representation of sensory data to estimate the input signals. Decoding experiments provide a quantitative characterization of information transmission and computational reliability under real-time conditions. The results of these studies show that neural coding and computation in several systems approach fundamental physical and informational theoretic limits to performance.

The world presents a rich variety of sensory signals, most of which may be thought of as varying continuously in time – sound pressure at the eardrum, light intensity in a region of the visual field, trajectories of objects, chemical concentrations, etc. Sensory data are carried to the brain not as continuously varying electrical signals, but as sequences of essentially identical discrete pulses, termed action potentials or spikes. The classic experiments of Adrian[1], Hartline (in Ref. 2), and co-workers identified many key features of this encoding. The rate of spike generation varies with the stimulus intensity (rate coding), the rate decreases in response to prolonged static stimulation (adaptation), and the rate can be affected by transformed versions of the stimulus, as in lateral inhibition (filtering or feature selectivity). Given the idea of rate coding, the strategy for experiments on sensory neurons is to plot the rate of firing as a function of a key parameter in the stimulus, mapping a 'receptive field'[3,4] or 'tuning curve'[5,6] for the cell in question. This approach has provided the basic language for thinking about sensory neurons for nearly 40 years; a significant proportion of what we know about the function and organization of the nervous system is based on measurements of neuronal firing rates.

Is the spike rate a complete description of the neural code? In the earliest experiments this rate was defined simply as the number of spikes in a given time interval following onset of the stimulus. In modern (computer assisted) experiments the rate is obtained by averaging multiple presentations of the same stimulus waveform, so that the rate defines the probability per unit time of a spike occurring. White noise or Wiener kernel methods eliminate the need for repeating a particular stimulus waveform many times, allowing measurement of stimulus–rate relations for an entire stimulus ensemble rather than just one waveform from the ensemble. These methods have been used extensively to study coding in the auditory and visual systems[7–11].

To go beyond rate coding, one can examine the statistical relations of the spikes to one another using interval distributions or correlation functions of the spike trains. This leads to classification of cells as being 'regular' or 'irregular', 'bursters', and so on. In several sensory systems one can find neurons with comparable firing rates but very different spike statistics (for example, Ref. 12), and in at least one case it was shown that a postsynaptic neuron gives different responses to presynaptic spike sequences with the same mean rate but different statistics[13]. These results raise the possibility that information is conveyed by more than just the rate, specifically by the timing of individual spikes. These ideas can be traced back at least to Wever's early work on synchronous activity in the auditory nerve[14]. Recent experiments on the mammalian visual system have again brought attention to the possible incompleteness of rate-based descriptions[15–17]. In addition to the possible importance of spike correlations as a coding scheme, many groups explored inter-neuron correlations as a diagnostic of network connectivity (see, for example, Ref. 18).

The best-studied example of a timing code is a code based on the distribution of inter-spike intervals[19,20]; the rate is the first moment of this distribution. Consider an auditory neuron stimulated at low frequencies ($\leq$1 kHz). The number of spikes in response to a tone-burst provides no information about phase, and there is a confusion between amplitude and frequency, since loud sounds away from the peak of the cell's frequency sensitivity produce the same firing rate as quiet sounds at the best frequency. But the phase-locking of auditory fibers causes the inter-spike intervals to cluster around integer multiples of the stimulus period[7]. This effect yields an independent estimate of the frequency, and the amplitude of the sound can then be estimated unambiguously from the firing rate. Thus by keeping track of spike arrival times, one can resolve the amplitude/frequency ambiguity that arises in a rate code.

Our interest in neural coding stems from a desire to answer two sets of questions about neurons as physical devices for computation and communication. Firstly, how much information can the spike train of a single neuron provide about continuous sensory inputs? Is there any way in which the representation of signals in spike trains constitutes an efficient representation? Secondly, how reliably can a given piece of the nervous system compute the quantities that control behavior? Does the coding of sensory signals in a spike format limit the kinds of computations which the nervous system can perform?

Different answers to these questions drive theoretical efforts in quite different directions. For example, if computation in small pieces of the nervous system is unreliable, the problem for the animal is how to synthesize reliable behavioral decisions out of a system with unreliable elements; this leads to sharp mathematical problems, as first posed by von Neumann[21]. If single cells are quite reliable, the problem is how to make maximal use of the available signal-to-noise ratio, which again leads us to well-posed mathematical problems[22,23]. As Bullock emphasized 20 years ago, measuring the reliability of computation in a spiking neuron requires that we make sense of its output; it is easy to confuse a complex code with noise[24]. Thus we are led back to the neural code.

Taking the organism's point of view

The conventional formulation of the neural coding problem misses an important point: firing rates, interval distributions and so on are average quantities, properties of an ensemble of spike trains rather than a single spike train. In the real world an animal may not have access to these averages. Consider again an auditory neuron. The distribution of spike intervals characterizes the response to a single tone of fixed amplitude and frequency. Real-world signals can be thought of as tones that are modulated (perhaps by large amounts) in both amplitude and frequency. If the modulations are slow compared to the mean interval, many spikes are fired before the parameters of the tone change significantly. In this case one can build up the interval distribution and thereby estimate the amplitude and frequency. However, modulations in many biologically significant sounds (speech, bat echolocation, frog calls, cricket chirps, etc.) occur on timescales comparable to the mean interval between spikes, so that sensory neurons can generate only a few spikes during the correlation time of the input signal. With such a small number of spikes we cannot accumulate a reasonable inter-spike interval distribution, nor even get a good estimate of the firing rate before the parameters of the stimulus change.

The fact that timescales in natural signals are comparable to typical inter-spike intervals is not confined to the auditory system. In the fly visual system, as discussed below, movements across the visual field result in the generation of a compensating torque within 30 ms, during which time the movement-sensitive neurons generate just a few spikes each. In the mammalian visual cortex, pre-attentively discriminable textures produce an average of just 1–3 spikes per cell within the 50–100 ms behavioral decision time[25], while optimally chosen moving gratings produce modulations of less than 3 spikes per 100 ms (see, for example, Ref. 26).

One might argue that larger spike counts are available by averaging over many cells. This is certainly not possible for the fly, which relies on only a handful of wide-field movement-sensitive neurons for stable flight. Even for other systems such averaging is an untested assumption; simple averaging throws away information carried by correlations in the spike trains of different cells. We now discuss how we can measure information transmission and reliability in individual cells without making assumptions about how information is coded.

We approach neural coding from the position of the organism and ask: 'Given the spike train defined by the set of arrival times $\{t_i\}$, what can we say about the unknown, time-varying stimulus waveform $s(t)$?'. This approach to neural coding is nearly opposite the conventional one, in which one seeks a model of the neuron's response to known external signals. Ideas much closer to our point of view were first advocated by FitzHugh[27] who used individual spike trains to make choices from a limited set of input stimuli. Barlow and Levick[28] carried out a similar analysis in experiments on detection and discrimination of intensity in the cat retina, and Barlow subsequently went on to investigate the statistical limits of neural processing for more complex tasks[29]. This work focused, however, on forced-choice discrimination between a small number of possible signals.

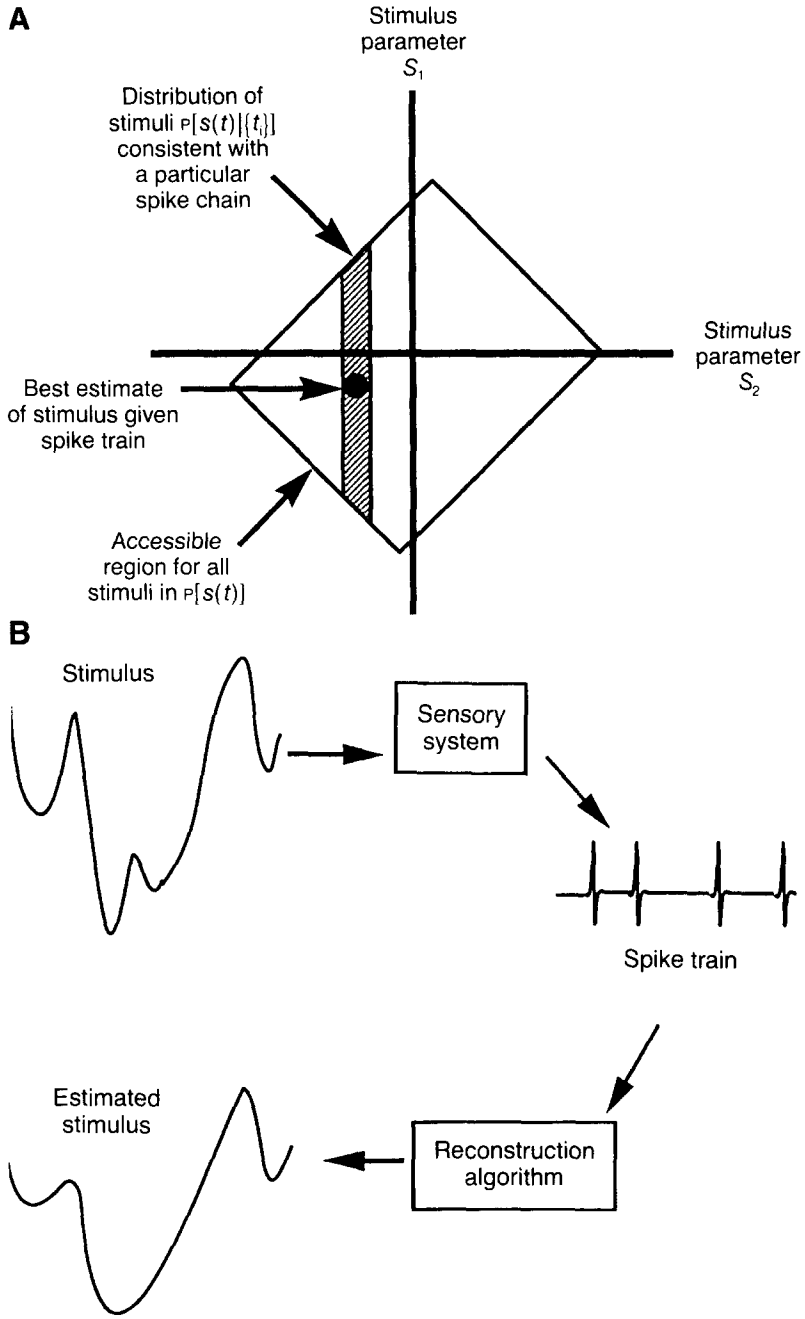

Fig. 1. (A) *Probability distributions. Stimulus waveforms are described by a large number of parameters (e.g. their Fourier coefficients). Here we show the distribution of two such parameters in the stimulus ensemble defined by* $P[s(t)]$ *(shown as box). For simplicity, we illustrate a case where the parameters are uniformly distributed over a region of the stimulus space. Once we observe the spike train, this distribution is narrowed* ($P[s(t)|\{t_i\}]$). *In this example, the spike train provides a great deal of information about one of the stimulus parameters but almost no information about the other. The quantitative measure of 'information' is the logarithm of the ratio of areas occupied by the two distributions. Our best estimate of the stimulus given the spike train is the centroid of the conditional distribution (shown by the filled circle).* **(B)** *A schematic of the reconstruction process. We place ourselves as observers of the spike train and devise an algorithm that takes the spikes as input and produces as output an estimate of the unknown, time-dependent stimulus waveform. The algorithm operates on a single example of the spike train and functions in real time so that the estimates are continuously updated. This algorithm is not the simple inverse of the input/output relation of the neuron.*

Johannesma, Gielen and Hesselmans developed a probabilistic formalism to address the more general problem of estimating an unknown, time-varying signal[30,31], and arrived independently at many of the ideas discussed in Refs 32, 33. In early applications

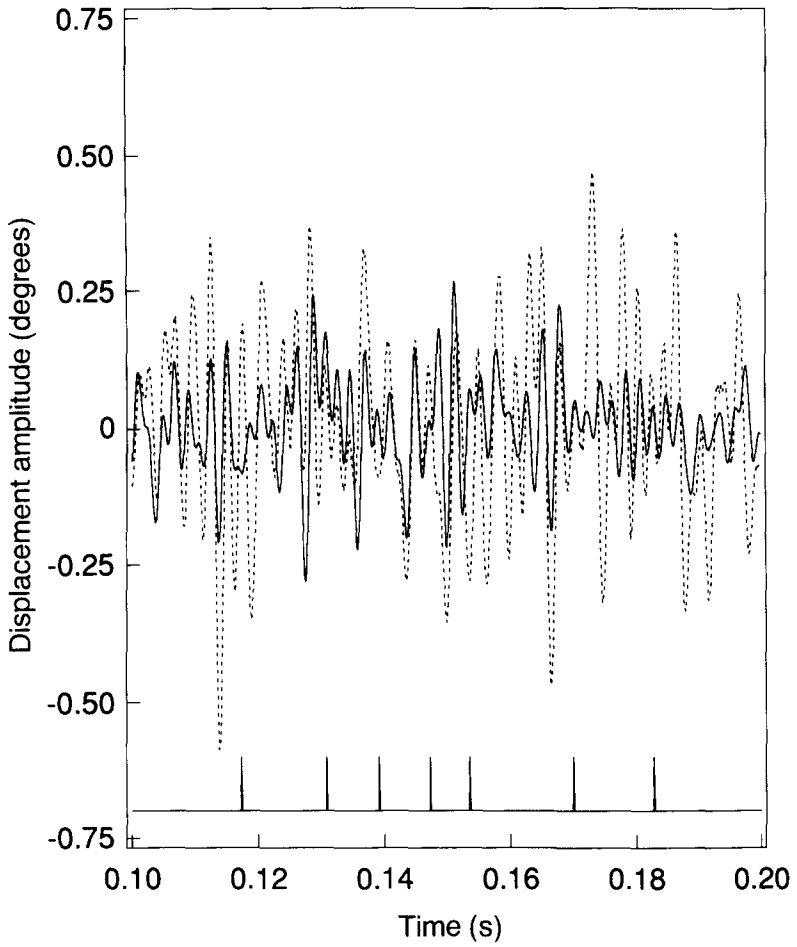

Fig. 2. *Signal (broken line) and reconstruction (unbroken line) in a single mechanosensor afferent from the cricket cercal system[39]. The stimulus consists of broad-band (25–525 Hz) Gaussian random motions of the sensory hair, and spikes (shown at the bottom) are recorded intracellularly. The reconstructions are accomplished by linear filtering of the spike train, as in Eqn 1, with the filter K_1 chosen to minimize the mean-square error between the signal and the reconstruction, as described in the text. Note that the spike sequence is extremely sparse, yet the reconstruction succeeds in capturing some of the high-frequency details in the waveform.*

of white-noise methods to the auditory system, de Boer[8] emphasized the interpretation of the reverse correlation function – the mean stimulus that triggers a spike – as the 'feature' of the stimulus waveform that is signalled by the occurrence of a spike (see also the review in Ref. 10). It was suggested[31] that one could estimate the stimulus waveform simply by adding up these features. This raises the problem that as the spikes arrive more frequently, the reverse correlation functions that are centered on successive spikes begin to overlap and can provide conflicting estimates of the stimulus. The correct resolution of these conflicts requires that we attach measures of confidence to the different estimates, and this is accomplished by measuring the relevant probability distributions[32], as described in the following section.

One of the most important aspects of the decoding approach is that the stimulus is unknown to the animal. In the natural environment or in an experiment, stimuli are chosen at random from some probability distribution $P[s(t)]$ that defines the 'stimulus ensemble'. Many experiments use simple ensembles (e.g. sine waves), but in these cases $s(t)$ can be predicted perfectly from knowledge of its past $s(t' < t)$. In this sense one can know the waveform without looking at the spike train, and there is no new information gained by observation of the spikes. Natural stimulus ensembles, on the other hand, are difficult to study because of their complex correlation structure. One might hope that a system could be characterized completely, so that its responses to signals drawn from any ensemble could be predicted. However, the nervous system is highly nonlinear so that the effective input/output relations on short timescales (receptive fields and tuning curves) depend on statistical features of the stimulus ensemble that are defined on longer timescales. We believe that it makes sense to take an empirical approach, studying the structure of the code in a variety of stimulus ensembles chosen to capture different aspects of the signals to which the animal is evolutionarily adapted.

Reconstructing the stimulus waveform

Assume that we record the spike times $\{t_i\}$ in a sensory neuron while stimulating the system with a continuous waveform chosen from a stimulus ensemble defined by the distribution $P[s(t)]$. With access only to the spike times, we, like the animal, do not know the stimulus. Before observing the spikes, all we know is that $s(t)$ was chosen from $P[s(t)]$; for any reasonably complex stimulus ensemble, this a priori knowledge gives only very crude information about the time-varying waveform. Once we observe the spike train, the set of possible waveforms narrows around the signals that are most likely to have generated that particular spike train, as schematized in Fig. 1A. This narrowing is described by the conditional probability distribution $P[s(t)|\{t_i\}]$, which measures the relative likelihood of different stimulus waveforms given a particular spike train. This distribution is everything that the animal could know about the stimulus as a result of observing the spike train.

The distribution of $P[s(t)|\{t_i\}]$ was measured for short spike sequences in experiments on a movement-sensitive cell (H1) of the fly visual system[32]. The structure of $P[s(t)|\{t_i\}]$ suggested that it should be possible to systematically reconstruct the waveform $s(t)$ using only the spike train $\{t_i\}$, and an attempt at this reconstruction was made[32]. The results of these studies suggested that we should try to decode the spike trains directly.

We describe the decoding process as a (generally nonlinear) filter operating on the spikes to produce an estimated stimulus:

$$s_{\text{est}}(t) = \sum_i K_1(t - t_i) + \tfrac{1}{2}\sum_{ij} K_2(t - t_i, t - t_j) + \ldots \qquad (1)$$

The filters K_n are related to the spiking statistics of the neuron and to the statistics of the stimulus ensemble (e.g. the power spectrum). In the context of a model of spike generation, the series is often dominated by the first (linear) term[33], so that the estimate takes the simple form of a linear filter applied to the spikes. The important point is that this series can be dominated by the linear term even when conventional measures of neural input/output relations are nonlinear. This means that it is in some sense easier to decode the spike train than it is to describe the encoding!

Eqn 1 describes a 'box' which could in principle be built out of real electronic (or neural) components. It

takes the spike train as an input and produces a continuous function of time that is our best estimate of the stimulus waveform; see Fig. 1B. When K_1 dominates, the box produces a stereotyped impulse response to each spike, and adds up the results. Timing of the spikes controls the estimate of $s(t)$ because the impulse responses to successive spikes overlap and interfere. In the limit that K_1 is very slow and the contributions of many spikes overlap, the output of the box resembles an estimate of the firing rate versus time; in this limit we recover a rate code which ignores fine variations in the times $\{t_i\}$. In the systems studied so far (see below), the impulse response K_1 typically overlaps just a few spikes, in accord with the discussion above.

Stimulus reconstruction is not necessarily a problem that is actually solved by the animal. It is, however, of the same nature as the problems that the animal must solve. As an example, the fly generates an opto-motor torque from the spike output of its motion-sensitive visual neurons, and this torque has a component roughly proportional to the time-dependent angular velocity[34,35]. The torque signal is a continuous analog waveform that the fly synthesizes out of discrete spike sequences in its sensory neurons. The fact that analog signals can be recovered so simply from the spike train is a fundamental characteristic of the neural representation of sensory data.

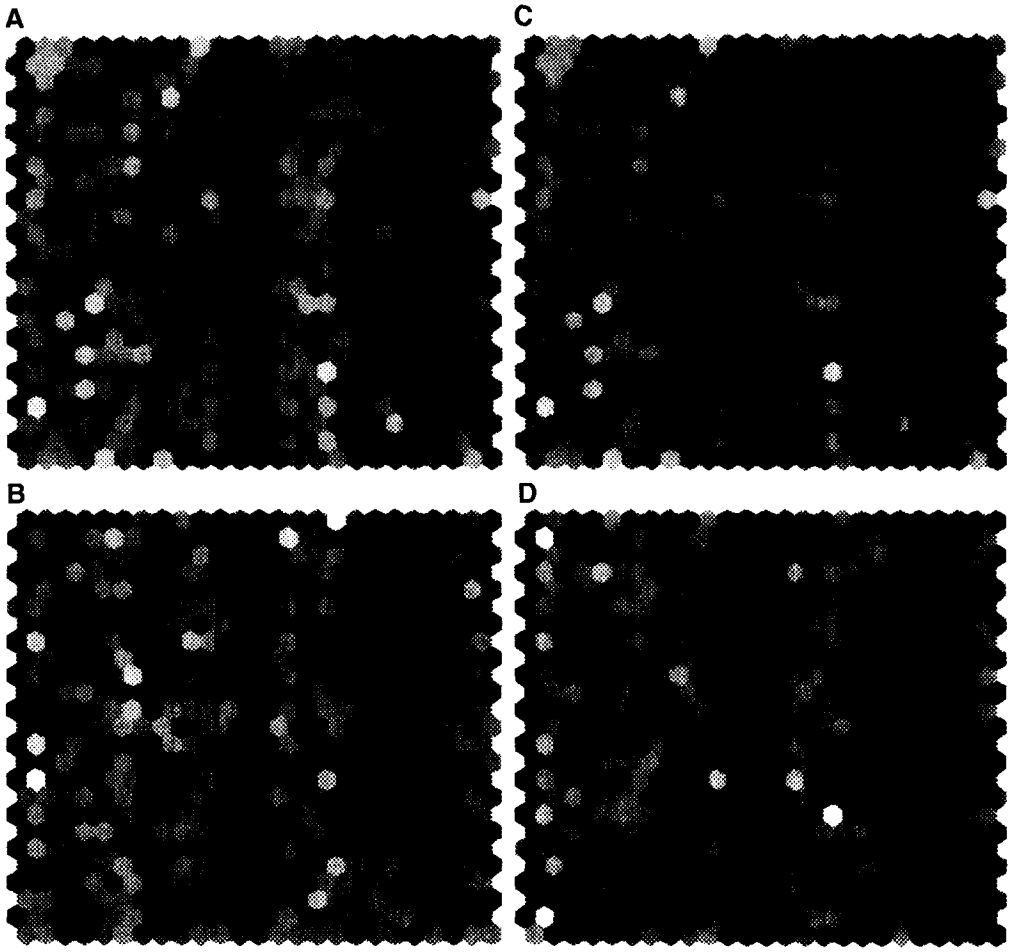

Fig. 3. *Snapshots of the fly photoreceptor voltage array in response to small displacements.* **(A), (B)** *Displacement of a random bar pattern by 0.36° [from* (A) *to* (B)*].* **(C), (D)** *Displacement of the same pattern by 0.48° [from* (C) *to* (D)*]. Simulations are based on signal and noise characteristics of the receptor cells measured under conditions identical to those used in the experiments on H1. In Ref. 29 it is shown that a single example of the spike train from H1 is sufficient to discriminate between these two displacements with a reliability ≥ 75%.*

Information rate and coding efficiency

When we observe the spike train we learn, in principle, about many different aspects of the stimulus. A frog call, for example, can be described by its fundamental frequency, the amplitudes and phases of the different harmonics, and the shape of the envelope. Each parameter can be estimated with a certain precision; measuring this precision in a single cell (see, for example, Refs 28, 36) is analogous to the psychophysical measurement of discrimination thresholds. However, discrimination thresholds measure performance when each stimulus dimension is isolated; in the real world, signals are varying continuously along all dimensions at once. How do we characterize the system's performance under these natural conditions?

Information theory provides a framework for adding up the many different discrimination abilities relevant to real-world signals[37,38]. In particular, the theory allows us to compare different stimulus dimensions (e.g. whether a given neuron provides more information about the amplitude or the pitch of a sound), even when discrimination abilities would be measured in different units (e.g. dB and Hz). We can think about the information in terms of the schematic in Fig. 1A. When we observe the spike train, the range of possible stimulus waveforms is narrowed into a smaller region of the stimulus space. The information provided by the spikes about the stimulus measures this reduction on a logarithmic scale, so that a factor of two reduction in the range of possible stimuli is counted as one bit of information. For example, imagine that frogs call with fundamental frequencies scattered uniformly throughout a 50 Hz range, and observation of the spike train of a single cell allows us to determine this frequency with a precision of 5 Hz. Then we gain $\log_2(50/5) \sim 3.3$ bits of information.

One way to measure the reduction in the range of possible stimuli given the spike train is to reconstruct the stimulus and measure the 'noise' in the reconstruction – the random errors of the reconstruction around the true stimulus. One can show mathematically that the information transmission rate can be estimated in terms of the variance or power spectrum of this noise, and that the true rate of information transmission is at least as high as our estimate. Together with David Warland, we used this approach to analyse experiments on primary afferent neurons from several different sensory systems[39–41] (see also Rieke, F. M. and Warland, D. K., PhD theses, University of California at Berkeley, 1991).

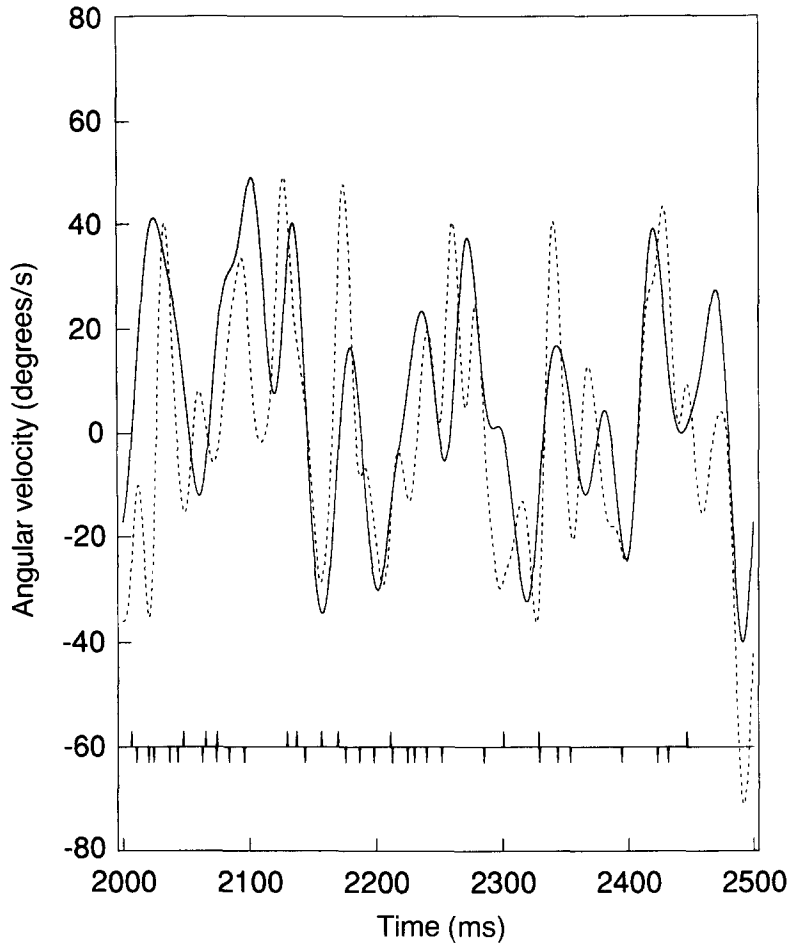

Fig. 4. *Time-dependent angular velocity signals (broken line) are reconstructed (unbroken line) by filtering the spike train of H1 (shown at the bottom). Positive spikes represent the response of H1 to the stimulus as shown; negative spikes represent the response to the negative of the stimulus, as would be seen by the H1 cell on the other side of the head during rigid rotation of the fly. The precision of the reconstructions can be quantified by computing the spectral density of the errors[22,46]. This measure of performance approaches the fundamental physical limits imposed by noise in the photoreceptors, as described in the text.*

Figure 2 illustrates one experiment on the mechanoreceptive neurons of the cricket cercal system. The stimulus is the angular displacement of the sensory hair, and spikes are recorded intracellularly from the primary afferent neuron. The reconstructed waveform clearly interpolates between the spikes, and in some places gives a close match to details of the stimulus on very short timescales. This tells us that the bandwidth of the system is large. On the other hand, the typical errors are comparable to the stimulus itself, so the overall signal-to-noise ratio is about one. Qualitatively, the estimate of the signal at any instant of time is imprecise, so that we obtain the order of one bit of information (the signal is positive or negative, for example). However, because of the large bandwidth, this estimate is updated very often, so the number of bits per second is large. Quantitative analysis shows that the single cell in Fig. 2 provides nearly 300 bits of information per second about the angular trajectory of the sensory hair. This is roughly three bits per spike (or per inter-spike interval), clear evidence that the timing of individual spikes carries significant information about the stimulus. Similar bit-to-spike ratios were obtained in experiments on neurons from the bullfrog sacculus, a vibratory sensory region in the inner ear[40].

In 1952, MacKay and McCulloch[42] pointed out that a system that keeps track of spike arrival times or inter-spike intervals could in principle convey several bits of information per spike, far more than a system relying only on firing rates averaged over a large time-window. Could the information rates we observe approach these limits? The absolute upper boundary to the transmitted information is set by the entropy of the spikes, which measures roughly the number of distinguishable spike sequences given some timing precision. The information rate measures the number of stimulus waveforms that can be distinguished from observation of the spikes. One cannot distinguish more waveforms than spike trains in any coding scheme. To compare the cell's performance to this fundamental limit, we estimated the spike-train entropy from the same experiments in which we found the high information rates. In both the cricket and the bullfrog, information rates are within a factor of two of the spike-train entropy, corresponding to a coding efficiency of greater than 50% (Ref. 41). This implies that although the spike sequences are highly variable, most of this variability is in fact being used to convey information about the stimulus.

Reliability in neural computation

To assess the reliability of neural computation we need to measure the signal and noise characteristics of the receptor cells providing the input data for the computation, and we need to measure the effective noise level of the signals carried by the neurons coding the output of the computation[22,23]. In the fly visual system, information about rigid rotation of the fly relative to the world is carried by a handful of identified wide-field neurons in the lobular complex; the input to these cells comes from a single class of photoreceptors in the compound eye[43,44]. Together with de Ruyter van Steveninck and Warland, we reported a series of experiments on the precision of the movement computation in the H1 cell, which senses rigid horizontal movements[32,36,45,46] (see also de Ruyter van Steveninck, R. R., Academisch Proefschrift, Rijksuniversiteit Groningen, 1986).

It turns out that the fly, like humans, can reliably distinguish displacements that differ by a small fraction of the spacing between receptors on the retina – the system achieves a displacement resolution much better than the nominal limit imposed by diffraction around the photoreceptor aperture. This class of phenomena is termed hyperacuity[47]. Microscopists have known for a century that one can resolve displacements far below the 'diffraction limit' provided that one has a sufficient signal-to-noise ratio. What is the signal-to-noise ratio in the fly's eye and how effectively is this signal-to-noise ratio used in displacement resolution?

Signal-to-noise ratios depend on the time over which we integrate the receptor response. In response to a sudden movement, the fly can generate a torque within 30 ms (Ref. 48). But most of this time is taken up in signal transmission from the retina to H1 (~ 15 ms), and in the rise time of the photoreceptor response (~ 10 ms). The remaining time is essentially equal to the correlation time in the photoreceptor voltage noise; so rather than looking at time-averaged

signals, the fly must base its behavioral decisions on one or perhaps two snapshots of the voltage array. The signal and noise characteristics of the photoreceptors were measured under conditions identical to the H1 recordings (de Ruyter van Steveninck, R. R., Academisch Proefschrift, Rijksuniversiteit Groningen, 1986), so we can simulate snapshots of the photoreceptor response to patterns that are displaced by small fractions of the receptor spacing, as in Fig. 3. This simulation demonstrates that on the timescales relevant to fly behavior the signals represented in the photoreceptor array are very noisy; and as a result, the hyperacuity task is hard. To quantify these results we developed the theory of the optimal motion estimator, an idealized device that uses the voltage signals from all of the eye's photoreceptors to generate the best possible estimate of the angular velocity waveform[22]. The measured photoreceptor noise limits the quality of this estimate, so that for movements in the frequency range from 10 to 25 Hz the minimum displacement noise is $\sim 0.1°$ in the behaviorally relevant 30 ms integration time. This is one-tenth of the angular spacing between receptors. The noise level of the optimal motion estimator sets an absolute standard against which we can judge the reliability of computation in the fly's visual system.

Following the approach described above, we decoded the spike train of H1 to estimate the trajectory of random patterns moving across the visual field; Fig. 4 shows an example of these reconstructions. The spike train of H1 contains enough information to infer – in real time, without averaging – details of the stimulus waveform on times comparable to the typical inter-spike interval, as suggested by the behavioral reaction times. Comparing these estimated waveforms with the real trajectories we found that the effective noise level in our reconstruction corresponds to an angular displacement noise within a factor of two of the theoretical limit imposed by the measured receptor cell noise. The discriminability of displacement differences in the range of $0.1°$ was confirmed directly in a separate series of experiments[36]. We conclude that the fly's visual system performs an optimal and nearly noiseless extraction of rigid movement signals from the photoreceptor array.

Looking ahead

By learning to decode neural spike trains we can quantitatively characterize the signals carried by sensory neurons. For the examples discussed here, this approach reveals that the nervous system reaches the fundamental physical and informational theoretic limits to reliability and efficiency. These examples fit into a growing body of data pointing to near optimal performance at different stages of sensory processing, as reviewed in Ref. 23. These data provide direct evidence that the theories of optimal coding and processing are relevant to the function and architecture of real neural circuits; in at least one case this argument can be carried to completion, resulting in successful parameter-free predictions of the signal transfer from photoreceptors to bipolar cells in the dark-adapted vertebrate retina[49,50]. To test these ideas fully we must study the way in which neural coding and computation are adapted to progressively more naturalistic ensembles of signals, and we must move away from the study of single cells to understand information transmission and coding efficiency for spike trains in arrays of neurons. Decoding methods have straightforward generalizations to these situations, and preliminary results are emerging[51,52].

Selected references

1 Adrian, E. D. (1928) *The Basis of Sensation,* W. W. Norton
2 Ratliff, F., ed. (1974) *Studies on Excitation and Inhibition in the Retina,* Rockefeller University Press
3 Kuffler, S. W. (1953) *J. Neurophys.* 16, 37–68
4 Barlow, H. B. (1953) *J. Physiol.* 119, 69–88
5 Galambos, R. and Davis, H. (1943) *J. Neurophys.* 6, 39–57
6 Kiang, N. Y-S., Watanabe, T., Thomas, E. C. and Clark, L. F. (1965) *Discharge Patterns of Single Fibers in the Cat's Auditory Nerve,* MIT Press
7 de Boer, E. and Kuyper, P. (1986) *IEEE Trans. Bio-Med. Eng.* 15, 169–179
8 de Boer, E. and de Jongh, H. R. (1978) *J. Acoust. Soc. Am.* 63, 115–135
9 Eggermont, J. J., Johannesma, P. I. M. and Aertsen, A. M. H. J. (1983) *Quart. Rev. Biophys.* 16, 341–414
10 Sakai, H., Naka, K-I. and Korenberg, M. (1988) *Vis. Neurosci.* 1, 287–296
11 Marmarelis, P. and Marmarelis, V. (1978) *Analysis of Physiological Systems. The White Noise Approach,* Plenum Press
12 Goldberg, J. M. and Fernandez, C. (1971) *J. Neurophys.* 34, 676–684
13 Segundo, J. P., Moore, G. P., Stensaas, L. J. and Bullock, T. H. (1963) *J. Exp. Biol.* 40, 643–667
14 Wever, E. G. (1949) *Theory of Hearing,* Wiley
15 McClurkin, J. W., Optican, L. M., Richmond, B. J. and Gawne, T. J. (1991) *Science* 253, 675–677
16 Gray, C. M. and Singer, W. (1989) *Proc. Natl Acad. Sci. USA* 86, 1698–1702
17 Eckhorn, R. *et al.* (1988) *Biol. Cybern.* 60, 121–130
18 Abeles, M. (1982) *Local Cortical Circuits: An Electrophysiological Study,* Springer-Verlag
19 Moore, G. P., Perkel, D. H. and Segundo, J. P. (1967) *Annu. Rev. Physiol.* 28, 493–522
20 Perkel, D. H. and Bullock, T. H. (1968) *Neurosci. Res. Prog. Bull.* 3, 405–427
21 von Neumann, J. (1956) in *Automata Studies* (Shannon, C. E. and McCarthy, J., eds), pp. 43–98, Princeton University Press
22 Bialek, W. (1990) in *1989 Lectures on Complex Systems, SFI Studies in the Sciences of Complexity* (Jen, E., ed.), pp. 513–595, Addison-Wesley
23 Bialek, W. (1992) *Analysis and Modeling of Neural Systems 2* (Eeckman, F., ed.), Kluwer Academic Press
24 Bullock, T. H. (1970) *J. Gen. Physiol.* 55, 656–684
25 Knierem, J. J. and van Essen, D. C. (1992) *J. Neurophys.* 67, 961–980
26 Reid, R. C., Sodak, R. E. and Shapley, R. M. (1991) *J. Neurophys.* 66, 505–529
27 FitzHugh, R. (1958) *J. Gen. Physiol.* 41, 675–692
28 Barlow, H. B. and Levick, W. (1969) *J. Physiol.* 200, 1–24
29 Barlow, H. B. (1980) *Philos. Trans. R. Soc. Lond. Ser. B* 290, 71–82
30 Johannesma, P. I. M. (1981) in *Advances in Physiological Sciences* (Szèkely, G., Lábos, E. and Damjanovich, S., eds), pp. 103–124, Akadémiai Kiadó
31 Gielen, C. C. A. M., Hesselmans, G. H. F. M. and Johannesma, P. I. M. (1988) *Math. Biosci.* 88, 15–35
32 de Ruyter van Steveninck, R. R. and Bialek, W. (1988) *Proc. R. Soc. Lond. Ser. B* 234, 379–414
33 Bialek, W. and Zee, A. (1990) *J. Stat. Phys.* 59, 103–115
34 Reichardt, W. and Poggio, T. (1976) *Quart. Rev. Biophys.* 9, 311–375
35 Poggio, T. and Reichardt, W. (1976) *Quart. Rev. Biophys.* 9, 377–438
36 de Ruyter van Steveninck, R. R. and Bialek, W. (1992) in *Advances in Neural Information Processing 4* (Lippmann, R., Moody, J. and Touretzky, D., eds), pp. 27–34, Morgan Kaufmann
37 Shannon, C. E. (1949) *Proc. IRE* 37, 10–21
38 Brillouin, L. (1962) *Science and Information Theory,* Academic Press
39 Warland, D., Landolfa, M., Miller, J. P. and Bialek, W. (1991) in *Analysis and Modeling of Neural Systems* (Eeckman, F., ed.), pp. 327–333, Kluwer Academic

Acknowledgements

The approach to neural coding advocated in this review was developed together with R. de Ruyter van Steveninck and D. Warland. We are indebted to many collaborators: D. A. Bodnar, R. R. Capranica, M. A. Landolfa, E. R. Lewis, M. Meister, J. P. Miller, D. Ruderman, W. Yamada and A. Zee. Finally, we thank the many friendly but critical colleagues who commented on these ideas in various stages of development: J. Atick, D. A. Baylor, H. B. Barlow, S. Block, T. H. Bullock, N. Franceschini, D. Glaser, A. Libchaber, E. Marder, W. G. Owen, J. Simmons and G. Zweig.

40 Rieke, F., Yamada, W., Moortgat, K., Lewis, E. R. and Bialek, W. (1992) in *Auditory Physiology and Perception: Proceedings of the 9th International Symposium on Hearing* (Cazals, Y., Demany, L. and Horner, K., eds), pp. 315–322, Elsevier
41 Rieke, F., Warland, D. and Bialek, W. (1992) in *Analysis and Modeling of Neural Systems 2* (Eeckman, F., ed.), Kluwer Academic
42 MacKay, D. and McCulloch, W. S. (1952) *Bull. Math. Biophys.* 14, 127–135
43 Hausen, K. (1984) in *Photoreception and Vision in Invertebrates* (Ali, M., ed.), pp. 523–559, Plenum Press
44 Hardie, R. C. and Stavenga, D. G., eds (1989) *Facets of Vision*, Springer-Verlag
45 de Ruyter van Steveninck, R. R., Bialek, W. and Zaagman, W. H. (1984) *Perception* 13, A47–A48
46 Bialek, W., Rieke, F., de Ruyter van Steveninck, R. R. and Warland, D. (1991) *Science* 252, 1854–1857
47 Westheimer, G. (1981) *Prog. Sens. Physiol.* 1, 1–30
48 Land, M. F. and Collett, T. S. (1974) *J. Comp. Physiol.* 89, 331–357
49 Bialek, W. and Owen, W. G. (1990) *Biophys. J.* 58, 1227–1233
50 Rieke, F., Owen, W. G. and Bialek, W. (1991) in *Advances in Neural Information Processing 3* (Lippman, R., Moody, J. and Touretzky, D., eds), pp. 377–383, Morgan Kaufmann
51 Bodnar, D., Rieke, F., Capranica, R. and Bialek, W. (1992) *Soc. Neurosci. Abstr.*
52 Warland, D. and Meister, M. (1992) *Soc. Neurosci. Abstr.* 18, 1191

Modelling of intersegmental coordination in the lamprey central pattern generator for locomotion

Avis H. Cohen, G. Bard Ermentrout, Tim Kiemel, Nancy Kopell, Karen A. Sigvardt and Thelma L. Williams

Avis H. Cohen is at the Dept of Zoology, University of Maryland, College Park, MD 20770, USA; G. Bard Ermentrout is at the Dept of Mathematics, University of Pittsburgh, Pittsburgh, PA 15260, USA; Tim Kiemel is at the Mathematical Research Branch, NIDDK, National Institutes of Health, Bethesda, MD 20982, USA; Nancy Kopell is at the Dept of Mathematics, Boston University, Boston, MA 02215, USA; Karen A. Sigvardt is at the Dept of Neurology, University of California-Davis, VAMC, Martinez, CA 94553, USA; and Thelma L. Williams is at the Physiology Dept, St George's Hospital Medical School, University of London, London, UK SW17 0RE.

Rhythmic motor activity requires coordination of different muscles or muscle groups so that they are all active with the same cycle duration and appropriate phase relationships. The neural mechanisms for such phase coupling in vertebrate locomotion are not known. Swimming in the lamprey is accomplished by the generation of a travelling wave of body curvature in which the phase coupling between segments is so controlled as to give approximately one full wavelength on the body at any swimming speed. This article reviews work that has combined mathematical analysis, biological experimentation and computer simulation to provide a conceptual framework within which intersegmental coordination can be investigated. Evidence is provided to suggest that in the lamprey, ascending coupling is dominant over descending coupling and controls the intersegmental phase lag during locomotion. The significance of long-range intersegmental coupling is also discussed.

Production of rhythmic motor patterns typically involves networks called central pattern generators (CPGs) that are believed to be composed of subnetworks of neurones, not necessarily identical, each capable of producing a rhythmic output; these are coordinated to provide proper relative timing. For example, in locomotion of limbed vertebrates, the network of neurons controlling the muscles of a single limb can be thought of as a group of such unit oscillators, each of which underlies the activity of a functionally related group of muscles. The mechanisms that produce the relative timing among the muscle groups have not been easily unravelled, largely because it is often difficult to understand or predict the behaviour of such oscillatory systems. In higher vertebrates such as mammals the problem is compounded by the large number of neurones involved in the respective CPGs.

Given these problems, the search for basic principles of oscillator coordination has focused on a more primitive vertebrate, the lamprey (reviewed in Refs 1, 2). Locomotor output patterns in the lamprey are conveniently simple: waves of lateral undulation travel down the body, propelling the animal through the water. Regardless of its speed, the lamprey maintains approximately one wavelength of curvature along its body at any time. The pattern of ventral root activity giving rise to this movement is correspondingly simple: there are alternating bursts of activity in the left and right motor roots of each segment, with a rostral–caudal delay of activation along the length of the spinal cord. Changes in speed are accomplished by either an increase or a decrease in the cycle duration at all segments with a proportional change in the intersegmental delay. Thus, the phase lag between segments (intersegmental time delay divided by cycle period; see Fig. 1A) is roughly constant at about 1% of the period per segment, independent of swimming frequency.

This phase delay cannot simply be composed of conduction delays or synaptic delays because these delays are constant, while the intersegmental delays must vary with the period to preserve the constant phase characteristic of the behaviour. It is intuitively difficult to see how a network might be connected to produce delays that remain a constant fraction of the cycle period, independent of frequency. Thus, it seemed prudent to turn to mathematical modelling for mechanisms that might govern the production of appropriate phase lags by coupled oscillators.

Since very little was known about the specific neurones involved in the local circuit oscillators or the intersegmental coordination, it was desirable to have a theory not tied to specific cellular mechanisms, but whose conclusions could nevertheless be tested experimentally. The theory described below[3–8] suggests mathematical mechanisms for production of the phase lags that are testable and which are compatible with a large range of underlying cellular mechanisms.

It is known that each segment or small group of segments from anywhere in the lamprey's body can produce bursting[9,10]. Thus, it is natural to describe the lamprey spinal cord as a chain of coupled limit cycle oscillators. A limit cycle oscillator is a dynamical system with a periodic solution to which the system returns after a perturbation, perhaps with a shift in