Physica A 200 (1993) 581–593
North-Holland

PHYSICA A

SDI: 0378-4371(93)E0148-8

Bits and brains: Information flow in the nervous system

William Bialek[a], Michael DeWeese[a,b], Fred Rieke[a,c] and David Warland[d]

[a]*NEC Research Institute, 4 Independence Way, Princeton, NJ 08540, USA*
[b]*Department of Physics, Princeton University, Princeton, NJ 08544, USA*
[c]*Committee on Neurobiology, University of Chicago, Chicago, IL 60637, USA*
[d]*Department of Cell and Development Biology, Harvard University, Cambridge, MA 02138, USA*

Until recently there have been no convincing quantitative measurements on the rates of information transmission in real neurons. Here we review the theoretical basis for making such measurements, together with the data which demonstrate remarkably high information rates in a variety of systems. In fact these rates are within a factor of two of the absolute physical limits set by the entropy of neural spike trains. These observations lead to sharp theoretical questions about the structure of the code and the strategy for adapting the code to different ensembles of input signals.

1. Introduction

Animals receive information from the world in the form of continuous functions of time – sound pressure at the eardrum, light intensity at each point in the visual field, the concentration of various substances in the air, ... At a very early stage in processing, however, these continuous signals are converted into discrete sequences of identical pulses, called action potentials or spikes. In each sensory neuron the spikes are all identical; information is carried only by the arrival times [1]. The dynamics by which these spikes are generated and propagated over relatively long distances from cell to cell are understood in molecular detail [10,25]; what is not known is how one should think about this encoding of the world from a computational or information theoretic point of view: How much information is carried by the spike train of a single cell? How does the representation of signals in spike trains affect the kinds of computations that the nervous system can accomplish? How is information shared among the many cells which, for example, carry signals from the eye or ear to the brain? Out of all the possible strategies for converting continuous signals into discrete spike trains, is there anything special about the strategy which is chosen in real brains? Is there anything universal about this strategy?

Recently it has become possible to give a much more quantitative analysis of neural coding in several different sensory systems. These experiments hint at some universal features of the code. The key experimental facts are described in two recent theses [19,30], and in several brief reports [6,20,22,31]; Ref. [21] is written for physicists. To summarize:

(i) *Very high information rates.* In the extreme, a mechanosensor in the cricket provides 300 bits per second of information about time-varying displacements of the sensory hair.

(ii) *High coding efficiency.* In the cricket and in a frog vibration sensor the rate at which the cell transmits information about the input signal is within a factor of two of the absolute limit set by the entropy per unit time of the spike train.

(iii) *Linear decoding.* Despite the evident non-linearities of the nervous system, spike trains can be *decoded* by simple linear filters. Specifically, if the input signal $s(t)$ is chosen from some reasonable probability distribution $P[s(\tau)]$ and the spikes arrive at times $\{t_i\}$, then the linear estimator

$$s_{\text{est}}(t) = \sum_i K_1(t - t_i) \tag{1}$$

is very good, and the addition of non-linear terms $\sim K_2(t - t_i, t - t_j)$ does not improve the quality of the estimate. This linear decoding exists under conditions where, for example, the relation between the input signal and the mean rate of spiking is strongly non-linear.

(iv) *Moderate signal-to-noise ratios.* All these examples of high information transmission rates correspond to cases where the estimate of the signal achieves a signal-to-noise ratio of order unity over a broad bandwidth, rather than a high SNR in a narrow band.

Several authors have argued that neural circuitry is in some sense designed to optimize the system's representation of the sensory world from an information theoretic point of view [2,3,7,12,13,24,27,29]. This is a very attractive hypothesis, and would fit with the general idea that the sensory systems reach fundamental physical limits to their performance [4,5]. Tests of the optimal coding idea, however, have been indirect – one tries to calculate some property of the optimal code (e.g., spatial filters or "receptive fields" for visual neurons) and check this against data on real neurons. But which features of the code are essential? Our strategy is to make a more direct attack on the problem, measuring information rates and comparing with the physical limits.

2. How we measure information rates

Until recently there were no direct measures of the rate at which the spike train $\{t_i\}$ in a sensory neuron provides information about continuously varying sensory signals $s(t)$. In the absence of data, speculations ranged over several orders of magnitude [14,28]. The information which the spike train provides about the stimulus is [9,26]

$$I[\{t_i\} \rightarrow s(\tau)] = \int \mathrm{D}t_i \int \mathrm{D}s\, P[s(\tau); \{t_i\}] \log_2\left(\frac{P[s(\tau); \{t_i\}]}{P[s(\tau)]\, P[\{t_i\}]}\right), \tag{2}$$

where $\int \mathrm{D}t_i$ is shorthand for integration over all arrival times $t_1, t_2, \ldots, t_N$ and summation over all spike counts N on the time interval $0 < t < T$, and $\int \mathrm{D}s$ as usual denotes an integration over the space of functions $s(t)$. Information is measured in bits, hence $\log_2$ $P[s(\tau)]$ is the a priori distribution from which the signal is drawn in a given experimental or natural situation. In the experiments of refs. [19,30] we start with the simplest case where this distribution is Gaussian and hence completely characterized by the signal power spectrum (two-point function) $S(\Omega)$. We can rewrite the transinformation in the form

$$\begin{aligned} I[\{t_i\} \rightarrow s(\tau)] = &- \int \mathrm{D}s\, P[s(\tau)] \log_2 P[s(\tau)] \\ &- \int \mathrm{D}t_i\, P[\{t_i\}] \left(- \int \mathrm{D}s\, P[s(\tau)|\{t_i\}] \log_2 P[s(\tau)|\{t_i\}]\right). \end{aligned} \tag{3}$$

The first term is the entropy of the distribution $P[s(t)]$, which we can assume is known since it is under the control of the experimenter. But the second term is the entropy of the *conditional* distribution $P[s(\tau)|\{t_i\}]$, averaged over the distribution of spike trains $P[\{t_i\}]$. Measuring the information transmission of a sensory neuron requires measuring this entropy of the conditional distribution, but in a real experiment all we have are samples drawn out of this distribution.

The problem of quantifying information transmission is rather like being asked to "measure" the entropy of a box of fluid. In fact this cannot be done, because the entropy is not an observable in the usual sense. One can measure entropy changes, but only because these are related to heat flows; there does not seem to be an informational analog of heat. As a compromise, one can *bound* the entropy of the fluid. At a trivial level, the entropy is surely less than that of an ideal gas of the same density. The accuracy of our estimate depends on how well we understand the correlations in the system.

The entropy of a distribution is always less than that of a Gaussian distribution having the same mean and variance. This leads us immediately to a lower bound on the transmitted information,

$$I \geq \frac{1}{2 \ln 2} \int \mathrm{D}t_i \, P[\{t_i\}] \operatorname{Tr} \ln[S\hat{N}^{-1}(\{t_i\})] , \tag{4}$$

where $\hat{N}(\{t_i\})$ is the covariance of the fluctuations in $s(t)$ around the conditional mean,

$$\bar{s}(t; \{t_i\}) \equiv \int \mathrm{D}s \, P[s(\tau)|\{t_i\}] \, s(t) , \tag{5}$$

and S is the covariance of the signal in the a priori distribution $P[s(\tau)]$. Finally since $\langle \ln x \rangle \leq \ln\langle x \rangle$, the rate of information transmission is

$$R_{\text{info}} = \lim_{T \to \infty} I[\{t_i\} \to s(\tau)]/T \geq \frac{1}{2 \ln 2} \int \frac{\mathrm{d}\Omega}{2\pi} \ln\left(\frac{S(\Omega)}{\bar{N}(\Omega)} \right) , \tag{6}$$

where $\bar{N}(\Omega)$ is the power spectrum of fluctuations around the conditional mean averaged over spike trains, and $S(\Omega)$ is the signal power spectrum as before. Note that although $\bar{N}(\{t_i\})$ described non-stationary fluctuations, the average over spike trains restores time-translation invariance and allows the definition of the power spectrum. If we construct some arbitrary estimator which takes as input the spike train $\{t_i\}$ and returns some estimate of $s(t)$, then the power spectrum of errors in this estimate $N_{\text{est}}(\Omega)$ will always be greater than or equal to $\bar{N}(\Omega)$, since the conditional mean is also the optimal least-square estimator. Hence

$$R_{\text{info}} \geq \frac{1}{2 \ln 2} \int \frac{\mathrm{d}\Omega}{2\pi} \ln\left(\frac{S(\Omega)}{N_{\text{est}}(\Omega)} \right) . \tag{7}$$

Clearly if we choose a bad estimator this bound will be far below the true information rate.

We thus arrive at a simple experimental strategy: Try to construct a box which takes as input the spike train $\{t_i\}$ and delivers as output an estimate $s_{\text{est}}(t)$ of the unknown, continuous stimulus chosen from the ensemble $P[s(\tau)]$. If we can parameterize this box, choose the parameters so as to minimize the mean-square deviation (χ^2) between the estimate and the true signal. Finally, the power spectrum of the errors will provide a lower bound on the information rate through eq. (7). In simple models [8] of the encoding process, the computation of the conditional mean $\bar{s}(t; \{t_i\})$ is analogous to the computation of the time-dependent expectation value of a quantum-mechanical coordinate subjected to an external force $F(t) \propto \Sigma_i \, \delta(t - t_i)$. This suggests writing our estimate in terms of the usual hierarchy of response functions,

$$s_{\text{est}}(t) = \sum_i K_1(t - t_i) + \frac{1}{2} \sum_{ij} K_2(t - t_i, t - t_j) + \cdots . \tag{8}$$

It is important to realize that this expansion is different than, for example, the expansion of time-dependent firing rates in terms of powers of the signal. Analysis of the model in ref. [8] shows that the input/output relation of the neuron can be strongly non-linear, yet s_{est} will be dominated by K_1 if the incoming signals are appropriately filtered; we return to this point below.

This approach was first tested in experiments on an identified neuron in the fly visual system [6]. The cell responds to rigid horizontal movements across the visual field, and produces a signal involved in the stabilization of straight flight. In this case the term K_1 is sufficient to give extremely precise estimates of the angular velocity waveform, and the addition of K_2 terms does not significantly reduce χ^2. Over a limited bandwidth we can make a stronger statement: The noise level in the estimates approach the limits imposed by noise in the photodetector array, so it would be impossible to improve the estimate no matter how many terms in the series (8) one wishes to include.

The success of linear decoding now extends to sensory neurons in a wide variety of systems: A simple mechanical sensor in the cricket [30,31], a vibration sensor in the frog inner ear [19,22], the two acoustic sensors of the frog inner ear [19,20], and the optic nerve of the salamander [32]. In each case "success" is first defined perturbatively to mean that inclusion of K_2 terms does not substantially reduce the errors in the estimated signal waveforms; in the one case where small but real improvements do occur, it was checked that K_3 does not produce statistically significant changes. In all these cases the errors in the reconstruction are Gaussian to a good approximation, so we expect that our bound in eq. (7) will be quite tight. The cricket provides the highest information rates, but the information per spike is comparable in the frog. In these cases we can also estimate the entropy of the spike train, and find that the information which the spikes provide about the stimulus is at least half of the spike train entropy. This shows that the coding efficiency is high – more than half of the degrees of freedom carry information – but also that the "success" of linear decoding is again not just perturbative, since it is not possible even in principle to extract much more information.

Finally, although the absolute noise levels in the reconstructed waveforms can be quite small – sufficient in the cricket to give nanometer precision in estimating the displacement of the sensory hair – the signal-to-noise ratio is moderate, of order one or less in many cases. If the Gaussian signals are embedded in a background of Gaussian noise, then from Shannon we know that the information rate is $R_{\mathrm{info}} = \Delta f \log_2(1 + \mathrm{SNR})$, where Δf is the bandwidth. Evidently widening the bandwidth provides a greater increase in information than a comparable increase in signal-to-noise. In the cricket the 300 bits per second in fact arises (roughly) from an SNR of unity over a 300 Hz bandwidth.

3. Why does linear decoding work?

We have presented the process of decoding spike trains purely as a device for estimating information transmission rates. However, it is remarkable that such a simple strategy as linear filtering produces estimates sufficient to extract several bits of information per spike. Why does this work? The key is to think about the correlation time τ_c of the signal (as seen through any filters in the neuron or previous cells) in relation to the typical inter-spike intervals. Clearly the occurrence of a single spike tells us something about the signal within a time window of roughly $\pm\tau_c$ around the spike itself. If the next spike occurs much later, with an interval $\tau \gg \tau_c$, then this spike gives us independent information about the signal and the contributions of the two spikes to our estimate of the stimulus waveform must just add. This suggests that our expansion of the optimal estimate in eq. (8) is really an expansion in $\langle r\rangle\tau_c$, and this can be verified by detailed perturbation theory calculations in the model of ref. [8].

Linear decodability thus defines a regime of neural dynamics in which each significant variation in the signal (on time scale τ_c) triggers of order one spike or less. This is almost the opposite picture from that suggested in rate coding models [28], where information is carried only in windows of time which contain several spikes, enough to form a reasonable estimate of the firing rate over the window. Is there any evidence concerning the value of $\langle r\rangle\tau_c$? Modulations in many biologically significant sounds (speech, bat echolocation, frog calls, cricket chirps, . . .) occur on time scales ~5–20 ms, during which time a cell firing 100 spikes per second can generate just one or two spikes. In the fly visual system movements across the visual field result in the generation of a compensating flight torque within 30 ms, during which time the handful of movement-sensitive neurons can generate just a few spikes each (see ref. [23]). In the mammalian visual cortex, preattentively discriminable textures produce an average of 1 to 3 spikes per cell within the 50–100 ms behavioral decision time [11], while optimally chosen moving gratings produce modulations of less than 3 spikes per 100 ms (see, for example, ref. [18]).

If we imagine that spike generation is a Poisson process whose rate $r(t)$ is determined by the stimulus $s(t)$, one can prove that the rate at which the spike train provides information about the stimulus is bounded by

$$R_{\text{info}} \leqslant \left\langle r(t) \log_2\left(\frac{r(t)}{\langle r \rangle}\right)\right\rangle . \tag{9}$$

The key to the proof is that the Poisson process has the maximum entropy of any point process with the same mean rate, and the inequality is saturated precisely in the limit $\langle r \rangle \tau_c \rightarrow 0$, which would guarantee linear decoding! This provides us with a hint that linear decoding may make sense for real neurons if the code has been "designed" to maximize information transmission.

4. Perturbation theory

As a model for spike generation let us imagine that the signals $s(t)$ are presented in a background of noise $\eta(t)$, and that the signal plus noise are passed through a filter $F(\tau)$. The resulting waveform

$$y(t) = \int \mathrm{d}r\, F(\tau)\, [s(t-\tau) + \eta(t-\tau)] \tag{10}$$

triggers a spike each time it crosses a threshold θ with positive slope. For simplicity let us assume that both the signal and noise are drawn from Gaussian distributions, characterized by the signal and noise spectral densities $S(\Omega)$ and $N(\Omega)$, respectively.

The threshold crossing model misses the refractoriness of real neurons, as well as the possibility that the cell is driven into a regime of self-sustained oscillation. Even without refractoriness, however, very short inter-spike intervals are suppressed. What we are missing should be negligible if the crossings are sufficiently infrequent, and in the extreme case the spikes will approximate a Poisson process. There is a substantial literature on the description of sensory neuron firing in terms of Poisson processes modulated by the incoming signal (see references in [5,8]). Although there are clearly deviations from this picture, it seems to be a good zero-order model of many systems. We will expand systematically around the Poisson limit, and try to show self-consistently that there is a near-Poisson optimal coding strategy. This leaves open the possibility that, in some cases, other optima exist far from the Poisson limit.

In the threshold crossing model the probability distribution of the spike train conditional on the signal is given by

$$P[\{t_i\}|s(t)] = \frac{1}{N!}\left\langle \exp\left(-\int \mathrm{d}r\, \rho(r)\right)\prod_{i=1}^{N} \rho(t_i)\right\rangle_\eta , \tag{11}$$

$$\rho(t) = \delta(y(t) - \theta)\, \dot{y}(t)\, H[\dot{y}(t)] , \tag{12}$$

where $\langle \cdots \rangle_\eta$ denotes an average over the distribution of the noise. We see that $P[\{t_i\}|s(t)]$ is the (normalized) average of a product of terms. In the Poisson limit each term is statistically independent, which suggests a cluster expansion,

$$\left\langle \prod_i g(t_i)\right\rangle = \prod_i \langle g(t_i)\rangle \left(1 + \frac{1}{2}\sum_{i\neq j} \frac{\langle \delta g(t_i)\, \delta g(t_j)\rangle}{\langle g(t_i)\rangle \langle g(t_j)\rangle} + \cdots\right). \tag{13}$$

If the correlation terms are small we can make use of the hypothesis that $\eta(t)$ is Gaussian to expand once more,

$$\begin{aligned}\langle g(t_i)\, g(t_j)\rangle_\eta &= \langle g(t_i)\rangle_\eta \langle g(t_j)\rangle_\eta \\ &\quad + \int \mathrm{d}t \int \mathrm{d}t'\, \langle \eta(t)\, \eta(t')\rangle \left\langle \frac{\delta g(t_i)}{\delta \eta(t)}\right\rangle_\eta \left\langle \frac{\delta g(t_j)}{\delta \eta(t')}\right\rangle_\eta + \cdots ,\end{aligned} \tag{14}$$

where of course

$$\langle \eta(t)\, \eta(t')\rangle = \int \frac{\mathrm{d}\Omega}{2\pi}\, \mathrm{e}^{-\mathrm{i}\Omega(t-t')} N(\Omega) . \tag{15}$$

These expressions are all we need to generate a systematic expansion of the information rate (M.D. and W.B., unpublished).

Motivated in part by the experimental results summarized above, let us look first at the small signal-to-noise ratio limit, where we can expand the distribution $P[\{t_i\}|s(t)]$ in powers of $s(t)$. In this limit the information transfer takes the simple form

$$R_{\mathrm{info}} = R_1 + R_2 + \cdots , \tag{16}$$

$$R_1 = \frac{r}{\sigma^2 \ln 2} \int \frac{\mathrm{d}\Omega}{2\pi}\, |\tilde{F}(\Omega)|^2 S(\Omega) \left[\ln\left(\frac{1}{2\pi r\tau}\right) + \frac{\pi}{4}(\Omega\tau)^2\right] , \tag{17}$$

$$R_2 = -\frac{2r^2}{\sigma^4 \ln 2} \int \frac{\mathrm{d}\Omega}{2\pi}\, |\tilde{F}(\Omega)|^4 S(\Omega)\, N(\Omega) \left[\ln\left(\frac{1}{2\pi r\tau}\right) + \frac{\pi}{4}(\Omega\tau)^2\right]^2 . \tag{18}$$

The first term arises from the Poisson limit, and the second term gives a correction proportional to the mean number of spikes per correlation time as seen at the output of the filter F. We have defined

$$\sigma^2 = \int \frac{d\Omega}{2\pi} |\tilde{F}(\Omega)|^2 N(\Omega) , \tag{19}$$

$$\tau^{-2} = \frac{1}{\sigma^2} \int \frac{d\Omega}{2\pi} |\tilde{F}(\Omega)|^2 \Omega^2 N(\Omega) , \tag{20}$$

and eliminated the threshold θ in favor of the mean spike rate

$$r = \frac{1}{2\pi\tau} \exp\left(-\frac{\theta^2}{2\sigma^2}\right) . \tag{21}$$

We now try to argue that this admittedly crude expansion contains the essence of the problem.

5. Optimal thresholds and adaptation

We can state a variational principle for the neural code, namely that the filter $F(\tau)$ and the threshold θ should be chosen so as to maximize R_{info}. Eqs. (16)–(18) give us an approximate expression for R_{info} which we can use to make this principle explicit. Clearly the optimal filters will depend on the signal and noise spectra, which makes sense. Let us suppose that these filters have been found. Then we are left with the problem of optimizing the threshold, or equivalently (from eq. (21)) the product $r\tau$.

First a question: What happens if the signal $s(t)$ is presented in a dc background, so that the mean signal is not zero (as assumed implicitly)? In a linear or weakly non-linear system, constant offset at the input results in constant offset at the input, with no effect on information transmission. In neurons, however, the only output of the cell which can be measured reliably is the arrival time of the spikes; all subthreshold voltage variations are attenuated in propagation along the axon until they are ultimately obscured by noise. As a result the output $\{t_i\}$ is not an invertible function of the input. Constant input signals are then relevant, and to maximize information transmission these dc terms *must* be ignored by the cell. Quantitatively, if the signal $s(t)$ acquires a dc offset, then the value of the threshold which maximizes the information transmission is offset by exactly the same amount, so that the constant signal has no effect on the firing rate.

It has been known since the initial experiments of Adrian, Hartline and coworkers that sensory neurons adapt to constant stimuli [1]. Traditionally it has been assumed that this adaptation represents a mechanism by which the organism chooses to ignore "uninteresting" static signals. In this point of view adaptation seems a sort of optional strategy; after all, why not wait for higher processing centers to decide what is truly "interesting"? The present analysis suggests that such an adaptation is necessary to insure maximum information transmission: The organism does not have the choice of discarding uninteresting signals at a later stage, since responding to the static signal would preclude the transmission of information about dynamic signals.

Returning to the problem of optimizing threshold, consider the simple case where both the signal and noise spectra are nearly constant (white), so that all of the dynamics are contributed by the filter F. To proceed analytically we want this filter to have a single characteristic time, so let us choose $|\tilde{F}(\Omega)|^2 \propto \exp[-(\Omega\tau)^2/2]$. Then the information rate becomes

$$R_{\text{info}} = \frac{S}{N}\frac{\tau^{-1}}{\ln 2}\, g(r\tau)\,, \tag{22}$$

$$g(x) = x\left[\ln\left(\frac{1}{2\pi x}\right) + \frac{\pi}{4}\right] - x^2(\pi)^{1/2}\left\{\left[\ln\left(\frac{1}{2\pi x}\right) + \frac{\pi}{8}\right]^2 + \frac{\pi^2}{32}\right\} + \cdots . \tag{23}$$

The optimum is reached when $r\tau = 0.12$, and at this point the correction to the information rate derived from inter-spike correlations is less than ten percent. This confirms that there exists an optimal setting of the threshold which results in a spike train near the Poisson limit. At the optimum $R_{\text{info}}/r = 1.58(S/N)$ bits/spike. If we extrapolate to $S/N = 1$ the prediction is within a factor of two of the observed 3 bits per spike observed in the cricket, although certainly higher terms in the S/N expansion should be significant at this point. In the frog acoustic sensors, where S/N in response to white signals is smaller, the information rates are close to the predicted ~ bit per spike.

6. Temporal jitter and coding efficiency

Until now we have assumed that the spike arrival times $\{t_i\}$ are measured with infinite precision. Imagine instead that the spikes are jittered with timing errors of standard deviation δr_{rms}. Then it can be shown that the first term in our expansion (R_1) is modified in a simple way,

$$R_1 \to \frac{r}{\sigma^2 \ln 2} \int \frac{\mathrm{d}\Omega}{2\pi} |\tilde{F}(\Omega)|^2 S(\Omega)\, [1 - (\Omega\, \delta\tau_{\mathrm{rms}})^2 + \cdots]$$
$$\times \left[\ln\left(\frac{1}{2\pi r\tau}\right) + \frac{\pi}{4} (\Omega\tau)^2 \right] . \tag{24}$$

If the spikes are counted in bins of width $\Delta\tau$, the same first order correction occurs but with $\delta\tau_{\mathrm{rms}} = \Delta\tau/\sqrt{12}$. In the case of the Gaussian filter discussed above,

$$R_1 \to R_1(\Delta\tau = 0)[1 - (3.56 r\, \Delta\tau)^2 + \cdots] , \tag{25}$$

where we have substituted the optimal relation between r and τ. Thus we see that ninety percent of the information is preserved when $r\,\Delta\tau \sim 0.09$, and ninety-five percent is preserved at $r\,\Delta\tau \sim 0.06$, within a factor of two of the results in the fly, cricket and frog vibration sensors [6,21]. We see that optimal information transmission is possible with only modest demands on timing precision.

When viewed through bins of (small) width $\Delta\tau$, a Poisson process has an entropy per event

$$S = -\frac{1}{r\,\Delta\tau} [r\,\Delta\tau \log_2(r\,\Delta\tau) + (1 - r\,\Delta\tau) \log_2(1 - r\,\Delta r)] , \tag{26}$$

or 4.85 bits per spike at $r\,\Delta\tau = 0.09$; in reality the entropy of the spike train will be slightly lower due to correlations, but these are small at the point of optimum information transmission. The coding efficiency is

$$\epsilon = \frac{R_{\mathrm{info}}(\Delta\tau)}{rS} = \frac{(0.9)1.58(S/N)}{4.85} \sim 0.3(S/N) . \tag{27}$$

Again this is within a factor of two of the best results, and in close agreement with cases where S/N is indeed small.

7. Tentative conclusions

The idea of coding with $\langle r \rangle \tau_{\mathrm{c}} \sim 1$ seems to tie together our different observations on the neural code. There is clearly an optimal setting of the spiking threshold at small $\langle r \rangle \tau_{\mathrm{c}}$, and information rates, coding efficiencies and fault tolerance of the code at this operating point are all in reasonable accord with experiment. If one carries the theory out to higher signal-to-noise ratios, the optimal setting of the threshold must adapt not only to added mean signals but also to the magnitude of the signal variance. This form of adaptation is not much discussed in the literature, although there is certainly evidence for it in the vertebrate retina [15]. Finally, everything we have said has a natural generalization to arrays of neurons, and it has recently become possible to explore these arrays experimentally [16,17].

Acknowledgements

We thank our many experimentalist colleagues who made this work possible: D. Bodnar, R. Capranica, M. Landolfa, J. Miller, E. Lewis, M. Meister, R. de Ruyter van Steveninck and W. Yamada. In addition we thank H. Barlow, A. Libchaber, M. Meister, M. Potters, D. Ruderman and R. de Ruyter van Steveninck for many helpful conversations.

References

[1] E.D. Adrian, The Basis of Sensation (Norton, New York, 1928).
[2] J.J. Atick, Could information theory provide an ecological theory of sensory processing?, in: Princeton Lectures on Biophysics, W. Bialek, ed. (World Scientific, Singapore) pp. 223–289.
[3] H.B. Barlow (1961), Possible principles underlying the transformation of sensory messages, in: Sensory Communication, W. Rosenblith, ed. (MIT Press, Cambridge, MA, 1961) pp. 217–234.
[4] W. Bialek, Ann. Rev. Biophys. Biophys. Chem. 16 (1987) 455.
[5] W. Bialek, Optimal signal processing in the nervous system, in: Princeton Lectures on Biophysics, W. Bialek, ed. (World Scientific, Singapore, 1992) pp. 321–401.
[6] W. Bialek, F. Rieke, R.R. de Ruyter van Steveninck and D. Warland, Science 252 (1991) 1854.
[7] W. Bialek, D.L. Ruderman and A. Zee, Optimal sampling of natural images: A design principle for the visual system?, in: Advances in Neural Information Processing, vol. 3, R.P. Lippman, J.E. Moody and D.S. Touretzky, eds. (Morgan Kaufmann, San Mateo, CA, 1991) pp. 363–369.
[8] W. Bialek and A. Zee, J. Stat. Phys. 59 (1990) 103.
[9] L. Brillouin, Science and Information Theory (Academic Press, New York, 1962).
[10] B. Katz, Nerve, Muscle, and Synapse (McGraw-Hill, New York, 1966).
[11] J.J. Knierem and D.C. van Essen, J. Neurophys. 67 (1992) 961.
[12] S.B. Laughlin, Z. Naturforsch. 36c (1981) 910.
[13] R. Linsker, An application of the principle of maximum information preservation to linear systems, in: Advances in Neural Information Processing, vol. 1, D. Touretzky, ed. (Morgan Kaufmann, San Mateo, CA, 1989) pp. 186–194.
[14] D. MacKay and W.S. McCulloch, Bull. Math. Biophys. 14 (1952) 127.
[15] M. Meister, private communication (1991).
[16] M. Meister, J. Pine and D.A. Baylor, Recording visual signals in the retina with a multielectrode array, preprint (1992).
[17] M. Meister, R.O.L. Wong, D.A. Baylor and C.J. Shatz, Science 252 (1991) 939.
[18] R.C. Reid, R.E. Sodak and R.M. Shapley, J. Neurophys. 66 (1991) 505.

[19] F. Rieke, Physical Principles Underlying Sensory Processing and Computation, Dissertation, University of California at Berkeley (1991).

[20] F. Rieke, D. Bodnar and W. Bialek, Coding of natural sound stimuli by the bullfrog auditory nerve: Phase, amplitude and information rates, in: Proc. Third Int. Congress of Neuroethology, Montreal (1992).

[21] F. Rieke, D. Warland and W. Bialek, Coding efficiency and information transmission in sensory neurons, Europhys. Lett. 22 (1993) 151.

[22] F. Rieke, W. Yamada, K. Moortgat, E.R. Lewis and W. Bialek, Real-time coding of complex signals in the auditory nerve, in: Auditory Physiology and Perception: Proc. 9th Int. Symp. on Hearing, Y. Cazals, L. Demany and K. Horner, eds. (Elsevier, Amsterdam, 1992) pp. 315–322.

[23] R. de Ruyter van Steveninck and W. Bialek, Proc. R. Soc. London Ser. B 234 (1988) 379.

[24] B. Sakitt and H.B. Barlow, Biol. Cybern. 43 (1982) 97.

[25] B. Sakmann and E. Neher, eds., Single Channel Recording (Plenum, New York, 1983).

[26] C.E. Shannon, Proc. IRE 37 (1949), 10.

[27] M.V. Srinivasan, S.B. Laughlin and A. Dubs, Proc. R. Soc. Lond. Ser. B 216 (1982) 427.

[28] R.B. Stein, Biophys. J. 7 (1967) 797.

[29] F. Theunissen and J.P. Miller, J. Neurophys. 66 (1991) 1690.

[30] D. Warland, Reading Between the Spikes: Real-Time Processing in Neural Systems, Dissertation, University of California at Berkeley (1991).

[31] D. Warland, M. Landolfa, J.P. Miller and W. Bialek, Reading between the spikes in the cercal filiform haor receptors of the cricket, in: Analysis and Modeling of Neural Systems, F. Eeckman, ed. (Kluwer, Dordrecht, 1991) pp. 327–333.

[32] D. Warland and M. Meister, unpublished.