Temporal filtering in retinal bipolar cells

Elements of an optimal computation?

William Bialek*‡ and W. Geoffrey Owen‡

Departments of *Physics and ‡Molecular and Cell Biology, University of California at Berkeley, Berkeley, California 94720 USA

ABSTRACT Recent experiments indicate that the dark-adapted vertebrate visual system can count photons with a reliability limited by dark noise in the rod photoreceptors themselves. This suggests that subsequent layers of the retina, responsible for signal processing, add little if any excess noise and extract all the available information. Given the signal and noise characteristics of the photoreceptors, what is the structure of such an optimal processor? We show that optimal estimates of time-varying light intensity can be accomplished by a two-stage filter, and we suggest that the first stage should be identified with the filtering which occurs at the first anatomical stage in retinal signal processing, signal transfer from the rod photoreceptor to the bipolar cell. This leads to parameter-free predictions of the bipolar cell response, which are in excellent agreement with experiments comparing rod and bipolar cell dynamics in the same retina. As far as we know this is the first case in which the computationally significant dynamics of a neuron could be predicted rather than modeled.

INTRODUCTION

The laws of physics impose fundamental limits on the performance of any device designed to sense the environment. Remarkably, a number of biological sensory systems reach these limits (1), the classic example being the ability of both vertebrate and invertebrate visual systems to count individual photons. Whereas it has long been recognized that this ability places important constraints on the mechanism of phototransduction (2), it has not been widely appreciated that single photon counting also requires efficient and reliable neural computation.

Many years ago Barlow (3) suggested that the reliability of photon counting in the behavioral response of a whole organism is limited by the rate at which the photopigment rhodopsin is spontaneously isomerized due to thermal noise. If correct, this hypothesis demands that the neural processing of single-photon signals be reliable in the extreme, because no significant noise must be added to the photoreceptor signal as it passes through the many layers of the visual system which contribute to stimulus detection and identification. Spontaneous isomerizations have been detected as a major noise source in recordings of the current flowing through a single-rod photoreceptor outer segment (4, 5). Comparison of these data with behavioral estimates of the "dark noise" in the human visual system lent strong support to Barlow's hypothesis, and recent measurements comparing the temperature dependencies of behavioral and physiological dark noise levels in the frog visual system leave little doubt about the essential correctness of this idea (6, 7).

These observations strongly suggest that the reliability of perception at low light levels is not limited by noise or by inefficiencies in the computational hardware of the visual system; rather the limit is set by noise in the primary sensory input itself (1). To insure that this is the case, signal processing and decision making processes in the retina and beyond must add at most a negligible amount of noise, and this processing must be such as to extract essentially all of the available information. Whereas more experiments are needed to confirm the suggestion that the visual system performs an optimal and nearly noiseless processing of single-photon signals, it seems profitable to explore the consequences of this hypothesis.

In this paper we present a theory for the design of a processor that takes as input the photocurrents or photovoltages from a collection of rod cells and gives as output an optimal estimate of the time-dependent light intensity. We find that this optimal processor can, at low photon fluxes, be naturally decomposed into two stages of filtering, and we suggest that these may be identified with the rod-bipolar and bipolar-ganglion cell signal transfer stages in real retinae. This being said, the dynamics of the hypothetical 'bipolar cell' response are completely determined by the signal and noise characteristics of the rods. Parameter-free predictions derived on this basis are in good agreement with recent experiments comparing rod and bipolar cell voltage responses in the same organism (8, 9). We discuss some implications of these results and suggest opportunities for more direct tests of our predictions.

Biophys. J. © Biophysical Society
Volume 58 November 1990 1227-1233

0006-3495/90/11/1227/07 $2.00

DESIGNING THE OPTIMAL FILTER

We pose the following problem: a collection of photoreceptors receives photons at a time-varying rate $R(t)$. This is the rate of a random Poisson process, with the photons themselves arriving at times t_μ; in principle we can also consider non-Poisson light sources, but under natural conditions we expect that dim sources will be Poisson. These arrivals are signaled by current pulses of stereotyped waveform $I_0(t - t_\mu)$, but these pulses are embedded in a background of noise $\delta I(t)$. We shall assume that the current noise is Gaussian, but much of the formal development can be generalized. Our task is to estimate the stimulus waveform $R(t)$ given an example of the current trace

$$I(t) = \sum_\mu I_0(t - t_\mu) + \delta I(t). \tag{1}$$

In what follows we give this estimation problem a probabilistic formulation, which allows us to use methods from statistical mechanics as a guide to the solution. The result, given in Eqs. 5 and 6, is that the stimulus $R(t)$ can be estimated by passing the current $I(t)$ through a linear filter. This conclusion is valid for small signals, and we argue that it is relevant to the processing of linear-range responses in the dark-adapted retina.

Everything that we know about the signal is contained in the conditional probability distribution $P[R(t)|I(t)]$. Note that because our random variables are functions of time this is a distribution functional, and all integrals over this distribution are functional integrals. To calculate this distribution we use Bayes' theorem,

$$P[R(t)|I(t)] = \frac{P[I(t)|R(t)]P[R(t)]}{P[I(t)]},$$

where $P[R(t)]$ is the *a priori* distribution defining the probability of different signals occurring in a given natural or experimental stimulus ensemble. We then introduce the photon arrival times t_μ, because

$$P[I(t)|R(t)] = \int Dt_\mu P[I(t)|\{t_\mu\}]P[\{t_\mu\}|R(t)],$$

where $\int Dt_\mu$ is shorthand for integration over all photon arrival times $t_1, t_2 \ldots, t_N$, and a sum over all photocounts N. The distribution of arrival times $P[\{t_\mu\}|R(t)]$ is the standard Poisson expression

$$P[\{t_\mu\}|R(t)]$$
$$= \frac{1}{N!}\exp\left[-\int d\tau R(\tau)\right]R(t_1)R(t_2)R(t_3)\cdots R(t_N).$$

Finally, if the current noise δI is Gaussian, we can write

$$P[I(t)|\{t_\mu\}] = Z^{-1}\exp\left\{-\frac{1}{2}\int dt \int dt' \cdot \left[I(t) - \sum_\mu I_0(t - t_\mu)\right]\int \frac{d\omega}{2\pi}\frac{e^{-i\omega(t-t')}}{S_I(\omega)} \cdot \left[I(t') - \sum_\nu I_0(t' - t_\nu)\right]\right\}, \tag{2}$$

with Z a normalization constant and $S_I(\omega)$ the current noise power spectrum.

Putting the pieces together we define the generating functional $W[J(t); I(t)]$ through

$$\exp\{W[J(\tau); I(\tau)]\} = \int DR \int Dt_\mu P[R(t)] \cdot P[\{t_\mu\}|R(t)P[I(t)|\{t_\mu\}]\exp\left[\int d\tau J(\tau)R(\tau)\right]. \tag{3}$$

It is then straightforward to show that the average stimulus waveform $R(t)$, given that we have seen the current signal $I(t)$, is

$$\langle R(t)\rangle_{I(t)} \equiv \int DR R(t) P[R(\tau)|I(\tau)] = \frac{\delta W[J(\tau); I(\tau)]}{\delta J(t)}\bigg|_{J=0}. \tag{4}$$

Because this approach is similar to the functional integral formulation of statistical mechanics or field theory (10), we have several calculational techniques at our disposal. In particular we are interested in the signal estimation problem at very low light levels, where the currents $I(t)$ themselves are expected to be small. Thus, we carry out a perturbation expansion in I to find

$$\langle R(t)\rangle_{I(t)} = \text{constants} + \int d\tau F(\tau)I(t - \tau) + \cdots, \tag{5}$$

$$F(\tau) = \int \frac{d\omega}{2\pi} e^{-i\omega\tau} S_R(\omega)\frac{\tilde{I}_0^*(\omega)}{S_I(\omega)}, \tag{6}$$

where $S_R(\omega)$ is the power spectrum of fluctuations in the photon arrival rate $R(t)$, $\tilde{I}_0(\omega)$ is the Fourier transform of the single photon current pulse, $\tilde{I}_0(\omega) = \int dte^{i\omega t}I_0(t)$, and $\tilde{I}_0^*(\omega)$ is its complex conjugate.

Eqs. 5 and 6 define the design of a filter that would operate on the photoreceptor currents and produce an optimal estimate of the photon arrival rate as a function of time. This result is the leading term of a perturbation series at small currents, which can in fact be viewed as a double expansion. One condition is that we are looking at the limit in which the optimal estimator is a linear filter. We feel that this is an appropriate limit both because there is a well-defined range of linear response for most cells in the retina, and because a realistic consideration of nonlinearity would require data on deviations from linearity in the photoreceptor response itself. Although many of the relevant experiments have been done (e.g., with different amplitude flashes), a complete characterization of nonlinear responses in the photon-counting regime does not exist. A second condition for the validity of our results is that the signal-to-noise ratio must be low. This is always the case at very low light levels, so that a dark-adapted retina is in fact adapted to processing of signals at low signal-to-noise ratio.

To summarize, Eqs. 5 and 6 determine the optimal estimator of photon arrival rate in the dark-adapted retina under conditions of linear response. We shall see that these ideas are in fact applicable to the understanding of retinal responses in this limit.

RELATING THE OPTIMAL FILTER TO RETINAL ARCHITECTURE

In the preceding section we have derived the optimal estimate of time-varying light intensities given that we observe the currents produced by an array of photoreceptors. Here we should point out that in the retina, second-order cells do not have access to the photocurrents, only to the photovoltages. We have chosen to formulate our theory in terms of currents, however, because the current signals and noise in individual rods are statistically independent. In contrast, because the rod cells are coupled, voltage noises of neighboring cells are correlated. To discuss the estimation of photon flux in diffuse stimuli we would need to know the voltage noise averaged over a number of receptors, or equivalently the low spatial frequency spectral density of the noise. Intracellular experiments, however, monitor the voltage noise at a single node in the rod network, which is an integral over all spatial frequencies. It is most convenient, then, to discuss the problem in terms of rod currents and imagine that our filters are phenomenological transimpedances between the rod cell and some hypothetical cell whose voltage represents an estimate of the time-varying light intensity.

We have chosen what is arguably the simplest of problems in computational vision: estimate the photon flux to a collection of uniformly illuminated receptors under conditions where the receptor responds linearly. Nonetheless this task is nontrivial, as expected from the classic literature on the recovery of linearly encoded signals in noise (11). Here we have two noise sources, one, the current fluctuations across the receptor cell membrane (which itself can be decomposed; see below), and the other, the random arrival of the photons at the receptor. It is interesting to note that our general approach to the estimation problem not only predicts the form of the linear filter which provides optimal estimates at low photon fluxes, it also defines (by extending the perturbation expansion) the conditions under which linear filtering ceases to be the optimal strategy.

The optimal filter $\tilde{F}(\omega)$ is naturally broken into two stages. One, $\tilde{I}_0^*(\omega)/S_I(\omega)$, is matched to the rod cell signal and noise characteristics independent of the stimulus characteristics. The other, $S_R(\omega)$, depends explicitly on stimulus ensemble. More generally, if we were to design a filter that estimates some more complex functional of the light intensity, it would turn out that the first stage of filtering remains necessary and unchanged, although the second stage of filtering will be different and in general nonlinear. The same first stage of filtering appears if we design, for example, the optimal estimator of rigid motions across the visual field (Rieke, F., D. Warland, R. R. de Ruyter van Steveninck, and W. Bialek, manuscript in preparation).

These considerations indicate that the first stage filter acts as a universal preprocessor of the rod output. In contrast, this first stage of filtering must be followed by different devices to deal with different computational problems. Therefore, we suggest that this first stage of filtering should be identified with the first anatomical stage in visual signal processing, the transmission of signals from the rods to the bipolar cells. Specifically, we propose that the output of the first stage filter be identified with the intracellular voltage response of the bipolar cell. This response is then completely determined by the filter characteristic $I_0^*(\omega)/S_I(\omega)$, and by the response of the rod cell itself, $I_0(\omega)$.

If we attempt to build the filter $I_0^*(\omega)/S_I(\omega)$, we immediately face a problem: in general, this filter is acausal and, hence, not physically realizable. Such acausality is a general feature of optimal estimation problems (11) and arises because the photoreceptor itself is causal, so that photon arrival at time t_0 influences the current $I(t > t_0)$. Conversely, observations of the current at time t inform us about photon arrivals at times *earlier* than t. This means that if we want to estimate the light intensity at a given instant of time, we would need access to the currents some time in the future, which is of course impossible.

Acausality means that we have to wait for some time to make the optimal estimate. More precisely, if observations of the current at time t contribute significantly to our estimate of the light intensity at time $t - \tau$, the optimal reconstruction of the intensity must lag the true signal by at least τ seconds. If the impulse response of the optimal filter is acausal (nonzero at negative times) but restricted to a compact time interval, then one can simply shift this impulse response to positive times by introducing a delay and smoothing out the behavior near $t = 0$.

If we examine the filter in the frequency domain then the shifting in time does not change the amplitude of the frequency response, only the phase. But a causal system of given amplitude characteristic has a minimum phase shift at each frequency (11); we must, thus, introduce enough delay to generate this minimum phase characteristic, and this in turn determines by how much our reconstruction will lag the true signal. If we introduce no phase shifts then we have an acausal filter that is physically unrealizable. If we introduce more than the minimum phase shift then the response of the filter is delayed by more than is necessary to insure causality: our estimates of the stimulus lag the real stimulus by an excessively long time.

In what follows we will construct optimal causal filters simply by finding the minimum phase filter with amplitude characteristic

$$|I_0^*(\omega)/S_I(\omega)|.$$

This filter has the smallest delay consistent with causality, so we pay the smallest penalty relative to the ideal acausal reconstruction. Further, minimum-phase filters have physically realizable inverse filters, which means that we are not throwing away any information in our filtering procedure.

We conclude this section with some discussion of the term "optimal." Clearly the definition of optimal estimate in the problem we posed at the beginning of this section is not unambiguous. We have chosen to discuss the device which estimates the average signal waveform, although we might also have considered the most likely waveform or some more complicated weighted average. Each of these estimates is optimal according to some definition; we might seek the minimum least-square error, the most reliable identification of brief transients, etc. Under some conditions reasonable variations in this definition of optimality do not result in large differences in the structure of the optimal estimate, as discussed recently in connection with the problem of decoding neural spike trains (12).

Even if we believe from the outset that the visual system performs an optimal computation, it may be difficult to discern the definition of optimality which has been forced upon the system by evolution; we may expect that different optimization principles are relevant for different organisms. In this first effort we have tried to adopt a simple and easily implemented version of the optimality hypothesis, within which the design of the optimal processor is especially straightforward. This means that we can easily compare the predicted processing algorithm with the characteristics of real cells.

COMPARISON WITH EXPERIMENT

To make these ideas concrete we make use of the data on rod photocurrents and noise in *Bufo marinus* taken by Baylor and co-workers (4). All of the cells exhibit very similar signal and noise properties, and in one particular cell the single-photon current pulse is of the form $I_0(t) = Ae^{-\alpha t}(1 - e^{-\alpha t})^3$, so that

$$\tilde{I}_0(\omega) = \frac{A/\alpha}{\Pi_{n=1}^4 (1 - i\omega/n\alpha)}. \qquad (7)$$

The current noise consists of two components, one of which corresponds to the spontaneous occurrence of discrete photonlike events and, hence, has a power spectrum proportional to $|\tilde{I}_0(\omega)|^2$. In a single cell this noise source is not Gaussian, but if we consider a bipolar cell which integrates over ~50 rods (8, 9) the Gaussian approximation is quite adequate. In addition there is a continuous component of the noise, which appears to be filtered through only two of the four time constants appearing in I_0. Quantitatively, therefore,

$$S_I(\omega) = \frac{S_c}{[1 + (\omega/2\alpha)^2][1 + (\omega/4\alpha)^2]} + \frac{S_d}{\Pi_{n=1}^4 [1 + (\omega/n\alpha)^2]}, \qquad (8)$$

with $S_c/S_d = 0.245$. The filter we would like to build is then

$$G(\omega) = \frac{I_0^*(\omega)}{S_I(\omega)} \propto \frac{\Pi_{n=1}^4 (1 - i\omega/n\alpha)}{1 + (S_c/S_d)[1 + (\omega/\alpha)^2][1 + (\omega/3\alpha)^2]}, \qquad (9)$$

which is acausal because it has poles appearing in conjugate pairs and, hence, in both halves of the complex ω plane. The corresponding causal, minimum-phase filter can be constructed analytically by converting each conjugate pair of poles into a double pole in the lower half plane, and we find

$$G(\omega) \rightarrow \frac{\Pi_{n=1}^4 (1 - i\omega/n\alpha)}{(\omega - \omega_0 + i\gamma)^2(\omega + \omega_0 + i\gamma)^2}, \qquad (10)$$

with $\omega_0 = 0.938\alpha$ and $\gamma = 2.425\alpha$. The response of our hypothetical bipolar cell is obtained by passing the photocurrent impulse response I_0 through this filter. Interestingly this serves simply to cancel the numerator in Eq. 10. Finally we transform back to the time domain to obtain the predicted impulse response of the bipolar cell,

$$V_{\text{bipolar}}(t > 0) \propto \exp(-\gamma t)[\sin(\omega_0 t) - \omega_0 t \cos(\omega_0 t)], \qquad (11)$$

where by causality $V_{\text{bipolar}}(t < 0) = 0$. Note that because $\gamma \gg \omega_0$ the oscillations expected from the sin and cos terms are essentially unobservable.

Ideally we would like to compare our predictions with recordings of bipolar cell responses in the *Bufo* retina under conditions as close as possible to the rod current experiments; alas such data are not available. In the tiger salamander *Ambystoma tigrinum* we have data on rod and bipolar-cell voltages taken from the same retina during the course of one experiment (8, 9). To make a meaningful comparison we, thus, convert the rod current impulse response of reference 5 into a voltage impulse response using the data and model of reference 13. In Fig 1 *a* we show the predicted bipolar cell response together with the rod voltage response.

We emphasize that the predicted bipolar cell dynamics is not a model of this cell. Given the measured characteristics of the rod cells, summarized by Eqs. 7 and 8 above, the bipolar cell response (Eq. 11) is completely determined by the principle of optimal computation. We draw attention to the two main features of the predicted bipolar cell responses, features which are echoed in the observed responses as shown in Fig. 1 *b*:

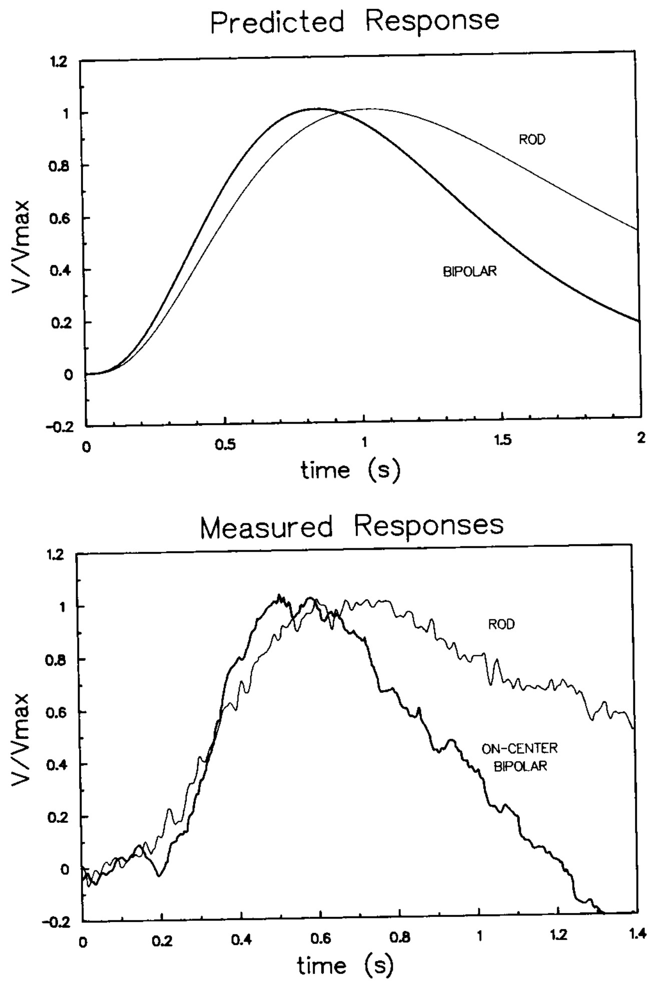

FIGURE 1 Theoretical and experimental responses in rod and bipolar cells. (*a*) Prediction of hypothetical bipolar cell voltage from Eq. 11 of the text, and rod voltage computed from data in references 5 and 13. (*b*) Measured voltage responses in a rod and an on-center bipolar cell from the tiger salamander (8, 9). Linear range responses are recorded to 20 ms flashes of 500 nm light, presented as circular spots 600 μm in diameter centered on each cell. Rod-cell data are taken with an intensity that bleaches 2 Rhodopsin per rod, bipolar cell data are taken at ~0.5 Rh*/rod. Each trace is the average of 10 intracellular records, which are then normalized; peak voltage excursions are 0.5 mV (rod) and 0.8 mV (bipolar).

(*a*) During the rising phase of the responses the bipolar cell response is predicted to lead the rod voltage, peaking roughly 10% earlier, in reasonable agreement with the experiment. The detailed prediction is, however, sensitive to the behavior of rod current signals and noise at high frequencies, where there is some variability from experiment to experiment (5). In addition, the short time behavior of the rod voltages in the salamander is significantly affected by the rod cell capacitance, which was neglected in the analysis of the rod network in *Bufo*. Despite these difficulties, it is clear that the basic prediction of a slightly quicker rise for the bipolar cell is confirmed.

(*b*) The bipolar cell voltage decays more rapidly than the rod cell voltage, so that the bipolar response is essentially complete whereas the rod response is still nearly half-maximal. Quantitatively, the full width at half maximum of the bipolar response is predicted to be ~35% less than that of the rod cell, and we observe ~40%. Similarly, the bipolar cell voltage falls to ≤25% of its maximum at times where the rod voltage is greater than half maximal. This quickening of the bipolar cell response relative to that of the rod means that the rod-bipolar transfer is at least in part a high-pass filter, as can be seen directly in the frequency domain (9). We recall that our theory is valid only for the case of linear responses in a dark-adapted retina. In fact the data of Fig. 1 *b* were taken in this limit (8, 9).

It is well known that rod voltage responses are faster than rod current responses, so that the rod network itself acts as a high-pass filter. We see from Fig. 1, however, that the extent of high-pass filtering required to convert the rod current into the bipolar cell voltage is still greater than that provided by the rods themselves. The apparent filter characteristics of the voltage-to-voltage transfer presumably reside in the rod-bipolar synapses, and the bipolar cell dendrites and soma.

In all this discussion we have been comparing salamander experiments with a theory that takes its parameters from data on a toad. This is dangerous, of course, but we know that the kinetics of the salamander photocurrents is similar in form to that of the toad (unpublished observations). A potentially more serious objection to our comparison is that rod cells seem to be significantly slower under the conditions of current recordings than under the conditions of voltage recordings, as may be seen by comparing references 4 and 13. This difficulty prevents us from placing Figs. 1 *a* and *b*, on the same absolute time scale, but the main predictions regarding the quickening of the bipolar response relative to that of the rod will not be affected. We could try to bridge this gap by treating the time scale α^{-1} (cf. Eq. 7) as a free parameter, but then we are reduced to curve fitting. If we allow ourselves this freedom we can achieve an almost perfect superposition of theoretical and experimental curves, as may be seen by carefully comparing Figs. 1, *a* and *b*. To give a decisive test of our predictions we must do careful measurements of rod voltage noise under the same conditions where we give a complete characterization of the rod network, because together these results can be used to infer the independent current signal and noise spectra. Such experiments are currently being analyzed (Rieke, F., and W. Bialek, manuscript in preparation).

DISCUSSION

In the last decade the subject of neural computation has blossomed into a large literature which crosses traditional boundaries among biology, physics, mathematics, and engineering. In those studies which concentrate on biology, one can discern two quite distinct traditions. One approach tries to motivate plausible models for the elements of neural computation, simplified neurons and synapses, and then studies the computational abilities which emerge when such elements are connected. Examples of this tradition can be found in the analysis of collective computation in highly interconnected networks (14, 15), and in the study of neural circuits which generate rhythmic behaviors (16, 17). An alternative approach is to focus on some computational task which the organism must face and then explore the minimal mechanisms that can extract signals of relevance to this task. An example of this approach is the study of correlation detection as a strategy for movement estimation (18, 19).

In this work we have suggested an approach to visual computation that we believe to be subtly different from these two traditions. In studying the processing of single-photon signals we have a very important piece of data, namely that this processing appears to be nearly optimal. We suggest that this optimality may be promoted to a design principle which can predict the functional dynamics of cells in the processing pathway. As a preliminary test of this principle we find the semiquantitative agreement between theory and experiment illustrated in Fig. 1 to be encouraging. As far as we know, there is no other case in which it has been possible to predict the dynamic response of a neuron from some design principle such as the idea of optimal computation proposed here.

We emphasize once more that our approach does not result in a model for the bipolar cell, but rather in a prediction of its response properties. As a result we can be wrong, and such errors cannot be fixed by curve-fitting. Failure of the optimality principle to correctly predict the computational abilities of cells in a processing pathway implies either that the postulated optimization principle is not relevant or that there are biological constraints which prevent optimization. Either of these possibilities is interesting, especially if one can quantify (as for photon counting) the approach to optimality in the system as a whole.

Our emphasis on optimal computation would be misplaced if the vertebrate retina provided an isolated example of this concept. In fact there is evidence of optimal processing in a number of visual tasks. In the fly retina it has been possible to compare effective contrast noise levels in photoreceptors and second-order cells, and one finds that at least over a limited frequency band no information is lost in this first stage of signal transfer (de Ruyter van Steveninck, R., personal communication). It is interesting that the temporal filtering which occurs between these two cells is similar to that found for the analogous rod bipolar signal transfer in vertebrates as discussed here. For the fly a picture of optimal processing would fit well with the optimality of photoreceptor optics in compound eyes (20).

Turning to more complex signal processing problems, recent evidence indicates that a movement-sensitive neuron in the fly visual system encodes all of the information available about rigid movements across the visual field, being limited only by the signal-to-noise ratio of the photoreceptors themselves (21–23). These observations provide a second opportunity to design the optimal processor and compare its dynamics with that observed for a real neuron. Finally, as emphasized by Barlow (24, 25), there is growing evidence that human observers can perform optimally or near-optimally at a number of perceptual tasks, including more complex *Gestalt* tasks such as recognition of symmetry. Perhaps theoretical methods along the lines presented here and in references 25 and 26 can be used to predict the character of neural computations which underlie such optimal performance.

We thank R. de Ruyter van Steveninck for doing the experiments which started us thinking about optimal computation and for sharing his observations on the fly retina in advance of publication. W. Hare and R. Miller helped collect and process the data of Fig. 1, and we had several helpful discussions with F. Rieke, D. Ruderman, V. Torre, and D. Warland.

This work was supported by the National Institutes of Health through Grant No. EY 03785 to WGO, and by the National Science Foundation through a Presidential Young Investigator Award to WB, supplemented by funds from Sun Microsystems and Cray Research.

Received for publication 31 October 1989 and in final form 15 June 1990.

Note added in proof: The in situ analysis of signals and noise in rods of the salamander retina has now been completed, allowing a fully quantitative and parameter-free test of the ideas presented here. The agreement between theory and experiment is excellent. A preliminary account will be given by Rieke, F., W. G. Owen, and W. Bialek. 1991. Optimal filtering in the salamander retina. *In* Analysis and Modeling of Neural Systems 1. F. H. Eeckman, editor. Kluwer Academic Publishers, Norwell, MA. An overview of these and related results is given in reference 28. A full account is in preparation.

REFERENCES

1. Bialek, W. 1987. Physical limits to sensation and perception. *Annu. Rev. Biophys. Biophys. Chem.* 16:455–478.
2. Hagins, W. 1965. Electrical signs of information flow in photoreceptors. *Cold Spring Harbor Symp. Quant. Biol.* 30:403–417.
3. Barlow, H. B. 1956. Retinal noise and absolute threshold. *J. Opt. Soc. Am.* 46:634–639.
4. Baylor, D. A., G. Matthews, and K.-W. Yau. 1980. Responses of retinal rods to single photons. *J. Physiol. (Lond.).* 309:591–621.
5. Baylor, D. A., B. J. Nunn, and J. F. Schnapf. 1984. The photocurrent and spectral sensitivity of rods of the monkey *Macaca fascicularis. J. Physiol. (Lond.).* 357:575–607.
6. Aho, A.-C., K. Donner, C. Hydén, L. O. Larsen, and T. Reuter. 1988. Low retinal noise in animals with low body temperature allows high visual sensitivity. *Nature (Lond.).* 324:348–350.
7. Barlow, H. B. 1988. The thermal limit to seeing. *Nature (Lond.).* 334:296.
8. Capovilla, M., W. A. Hare, and W. G. Owen. 1987. Voltage gain of signal transfer from retinal rods to bipolar cells in the tiger salamander. *J. Physiol. (Lond.).* 391:125–140.
9. Owen, W. G., W. A. Hare, and A. Wang. 1989. The role of the bipolar cell in retinal signal processing. *In* Proceedings of 1989 IEEE International Conference on Systems, Man, and Cybernetics. 435–442.
10. Parisi, G. 1988. Statistical Field Theory. Addison-Wesley Publishing Co., Reading, MA.
11. Wiener, N. 1949. Extrapolation, Interpolation and Smoothing of Time Series. John Wiley & Sons, New York.
12. Bialek, W., and A. Zee. 1990. Coding and computation with neural spike trains. *J. Stat. Phys.* 59:103–115.
13. Torre, V., and W. G. Owen. 1983. High-pass filtering of small signals by the rod network in the retina of the toad, *Bufo marinus. Biophys. J.* 41:305–324.
14. Hopfield, J. J. 1982. Neural networks and physical systems with emergent computational abilities. *Proc. Natl. Acad. Sci. USA.* 79:2554–2558.
15. Hopfield, J. J., and D. Tank. 1986. Computing with neural circuits. A model. *Science (Wash. DC).* 233:625–633.
16. Koppel, N. 1986. Oscillations in Chemistry and Biology. H. Othmer, editor. Springer-Verlag, Berlin.
17. Kleinfeld, D., and H. Sompolinsky. 1988. Associative neural network model for the generation of temporal patterns. Theory and application to central pattern generators. *Biophys. J.* 54:1039–1051.
18. Reichardt, W., and T. Poggio. 1976. Visual control of orientation behavior in the fly. I: a quantitative analysis. *Q. Rev. Biophys.* 9:311–376.
19. Poggio, T., and W. Reichardt. 1981. *In* Theoretical Approaches in Neurobiology. W. Reichardt and T. Poggio, editors. MIT Press, Cambridge, MA.
20. Snyder, A. W., D. S. Stavenga, and S. B. Laughlin. 1977. Information capacity of compound eyes. *J. Comp. Physiol.* 116:183–207.
21. Bialek, W., R. R. de Ruyter van Steveninck, and W. H. Zaagman. 1986. Reliability and information capacity of a neural code. Institute for Theoretical Physics Report NSF-ITP-86-117. University of California, Santa Barbara.
22. de Ruyter van Steveninck, R., and W. Bialek. 1988. Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proc. R. Soc. Lond. B Biol.* 234:379–414.
23. Bialek, W., F. Rieke, R. R. de Ruyter van Steveninck, and D. Warland. 1990. Reading a neural code. *In* Advances in Neural Information Processing 2. D. Touretzky, editor. Morgan Kauffman. 36–43.
24. Barlow, H. B. 1980. The absolute efficiency of perceptual decisions. *Phil. Trans. R. Soc. Lond. B Biol.* 290:71–82.
25. Barlow, H. B. 1981. The Ferrier Lecture, 1980: critical limiting factors in the design of the eye and visual cortex. *Proc. R. Soc. Lond. B Biol.* 212:1–34.
26. Bialek, W., and A. Zee. 1987. Statistical mechanics and invariant perception. *Phys. Rev. Lett.* 58:741–744.
27. Bialek, W., and A. Zee. 1988. Understanding the efficiency of human perception. *Phys. Rev. Lett.* 61:1512–1515.
28. Bialek, W. 1990. Theoretical physics meets experimental neurobiology. *In* 1989 Lectures in Complex Systems, SFI Studies in the Sciences of Complexity. E. Jen, editor. Addison-Wesley Co., Reading, MA. 513–595.