

The Neural Code of the Retina

Review

Markus Meister* and Michael J. Berry II
Department of Molecular and Cellular Biology
Harvard University
Cambridge, Massachusetts 02138

Introduction

How Do Action Potentials Represent Information?

Action potentials are the standard signal conveyed between neurons in the central nervous system. It is a long-standing question how these spikes represent sensory input, internal states of the brain, or motor commands (Perkel and Bullock, 1968; Rieke et al., 1997). To fully understand communication among neurons, one would like to obtain a dictionary for this language, in which each spike or pattern of spikes is assigned meaning within the processing task under study. This review will focus on the neural code employed by the ganglion cells of the vertebrate retina in conveying visual information from the eye to the brain: what are the rules by which the spike trains of optic nerve fibers encode the visual scene?

Communication between the retina and the brain is particularly amenable to experimental analysis for several reasons. First, we know exactly what is being represented by these action potentials, namely the time-dependent visual image as projected by the optics of the eye. Second, one can readily stimulate the retina with its natural sensory input, making use of well developed technology for presenting movies. Similarly, the output of the retina can be monitored with relative ease by extracellular recording from ganglion cells, optic nerve fibers, or terminals in the lateral geniculate nucleus. Finally, our retina performs a significant amount of processing, which compresses the visual signal from a neural population of 10^8 photoreceptors into just 10^6 optic nerve fibers. Nowhere else in the visual system is the scene represented with as few neurons as in the optic nerve, and thus one might expect to discover interesting principles of efficient coding. Since the electrical spikes on ganglion cell axons are the only source of our visual experience, there is considerable interest in the power of retinal processing and how it shapes our visual perception.

Our understanding of retinal coding has come a long way since the pioneering recordings from retinal ganglion cells by Kuffler (1953) and Barlow (1953). Sadly, even recent Neuroscience textbooks limit themselves to a qualitative treatment that does not reflect these advances. One goal of this review is to illustrate how the relationship between visual images and optic nerve firings can be captured quantitatively, to a degree that comes close to the ideal of the “dictionary” mentioned above. We then highlight some recent observations that have drawn attention to novel aspects of retinal processing.

A Sample Problem

What should a useful description of the retinal code contain? One way to find out is to try and answer a simple question, such as: “How many levels of gray can the retina distinguish?” This problem is taken from real life: it was posed to us by a colleague who expected a quick answer. It is also of some practical relevance, given that the price of home video systems is strongly related to how many gray levels they will produce. To tackle it experimentally, we imagine taking the following approach: record the spike train of a retinal ganglion cell, project a uniform gray field on the retina, vary the intensity of the light in small step increments, and ask how small the steps can be to still cause a recognizable difference in neural firing.

We soon find that after an intensity step the ganglion cell fires a brief burst of spikes, then settles down within a few seconds to whatever it was doing before the step. Something similar happens for almost every other ganglion cell, so we conclude that steady gray levels are almost indistinguishable. To get a meaningful answer, we therefore decide to vary the light intensity in time, for example switching back and forth between two levels. As we switch faster and faster, we find that the neural response eventually disappears again at very high frequencies. So the initial question regarding resolution of gray levels cannot be answered without specifying the time course of the light intensity, since there is no response for either very slow or very fast changes. Similarly, we soon find that the spatial distribution of intensity on the retina is very important. For example, when we illuminate only a small spot overlying the recorded ganglion cell, we find that the response is stronger than when we illuminate the entire field uniformly. On the other hand, the spot cannot be too small, or the response disappears again.

Next, we need to decide what it means to “distinguish” two gray levels. Presumably, someone monitoring the ganglion cell’s spike train should be able to identify which gray level is being presented with some degree of confidence. We find that even when presenting the same stimulus repeatedly, the ganglion cell produces somewhat different spike trains, and this variability ultimately limits the ability to discriminate two different stimuli. So, in seeking an answer to the initial question, the variability of the neural response is equally important as the response itself.

To quantitate any of this, we need to decide what aspects of the spike train to measure as the “response.” Is it sufficient to simply count the number of spikes in some suitable time window, or should we note the exact time of arrival of every spike? Furthermore, we must consider that many ganglion cells are being affected by this stimulus and can therefore contribute to its identification. So, we should try to understand how the responses of different neurons interact.

Finally, toward the end of the day, we discover a vexing feature of this experiment: if we gradually step the light intensity upward, we get one answer, and if we gradually step it downward, we get another. Apparently,

* To whom correspondence should be addressed (e-mail: meister@biosun.harvard.edu).

the relationship between gray level and firing is not permanent and static but varies considerably depending on the recent history of the visual stimulus.

This thought experiment illustrates that to answer a seemingly straightforward question about retinal signaling, one needs to know many different facets of how the visual information is encoded. In particular, any useful description of the neural code should specify: (1) the relevant measure of neural activity in the ganglion cell population; (2) how this activity responds to any given visual stimulus; (3) the precision of this response; and (4) the degree of plasticity in this relationship between stimulus and response, specifically, how it varies depending on recent history of the visual input.

Of course, similar issues arise in every study of neural communication, whether it regards sensory encoding, signaling among neural populations, or motor control.

The Prevailing View of Retinal Signaling

In this section, we attempt to summarize a consensus notion of how the retina encodes visual stimuli, delineating the four essential components of the neural code introduced above. This understanding has been gained from experiments on a wide range of species, and a comment is in order on how these should be integrated. Different animals clearly employ their visual system for different tasks, and this is reflected in the anatomical and functional diversity of their visual pathways (Stone, 1983). However, while the postretinal anatomy differs significantly across vertebrates, the structure of the retina is remarkably conserved from fish to primates. One finds the same three-layered arrangement, the same five principal cell types, the same neurotransmitters employed, and in many cases, the same anatomical microcircuitry (Dowling, 1987). A plausible explanation is that the retina is adapted to deal with constraints that are shared among all species: the statistics of visual images from the natural world at one end, and the limited capacity of the optic nerve at the other end. At any rate, many principles of retinal signaling seem to be remarkably conserved. The basic aspects of spatiotemporal processing, light adaptation, contrast gain control, and stochastic variation of the response are documented in animals ranging from tiger salamander to macaque monkey. Models of the light response that successfully predict a ganglion cell's firing rate share a common structure in all these cases. Differences among species affect the quantitative parameters of these models, but not their basic elements.

The Relevant Features of Ganglion Cell Spike Trains

Although the visual scene is conveyed to the brain in parallel by the spike trains of all optic nerve fibers, most of what we know about retinal signaling is derived from recordings of single retinal ganglion cells, one at a time. The same is true, of course, for neural signaling everywhere else in the nervous system. The underlying and often unstated assumption is that such a population code can, in fact, be understood one cell at a time. Two conditions are necessary for this. First, there should be identifiable classes of cells, that group neurons of similar functional properties, such that studying one or a few cells of a given class allows one to estimate the behavior

of other cells in this class. If, instead, the population were perfectly heterogeneous, the code could only be understood after observing every neuron. A great deal of effort has gone into sorting retinal ganglion cells into different types based on their visual responses. There is clear evidence for distinct classes—some of which will be discussed below—though their precise number and boundaries are often in dispute (for review see Stone, 1983; Rodieck et al., 1993; Rodieck, 1998). Second, the firing of each neuron in the population should depend only on the stimulus, not on the activity of other neurons in the population. If this is so, then the description derived from many single-cell recordings can adequately predict the occurrence of any given response pattern in the population. Unfortunately, this latter condition is not always met, as demonstrated in recent multineuron recordings discussed below. Nevertheless, the classical single-neuron analysis of retinal responses has been highly successful and valuable and is still the focus of much research.

Within the spike train of a single ganglion cell, the important response feature is generally taken to be the neuron's instantaneous firing probability at various times throughout the stimulus presentation. In experiments, this firing rate is estimated by repeating the same stimulus many times and counting spikes in the corresponding time bin of many such trials. Of course, during natural vision, we do not enjoy the luxury of many identical stimulus trials. It is often assumed that the brain, instead, estimates this response function by counting spikes from many essentially identical ganglion cells (Levick and Zacks, 1970; Enroth-Cugell et al., 1983). This idea conflicts with another commonly held notion, namely that ganglion cells of any given functional type "tile" the retina, such that each point is serviced by just one neuron of that type (Wässle and Boycott, 1991; DeVries and Baylor, 1997). We will revisit this topic below when considering distributed coding by retinal ganglion cells.

The Stimulus-Response Relationship

With these assumptions, the central problem of the retinal code is how a ganglion cell's firing rate depends on visual stimulation. Early experiments explored this relationship with simple stimuli, such as a small spot flashed on a uniform background (Barlow, 1953; Kuffler, 1953). Generally, the spot altered the firing rate only if presented within a small region on the retina—termed the "receptive field"—a few tens to hundreds of micrometers in diameter surrounding the cell body. The nature of the light response within the receptive field immediately pointed to the existence of very different cell types. In some ganglion cells, a spot flashed near the center of the receptive field produced a transient increase of firing at light onset and a brief reduction of firing at offset (ON cells). In other ganglion cells, the firing rate decreased at onset and increased at offset (OFF cells). For either cell type, a spot placed at some distance from the center—in the so-called receptive field surround—had the opposite effect of a spot in the center. When center and surround were illuminated simultaneously, the center response was significantly suppressed. Still other ganglion cells responded with a brief burst of spikes at both onset and offset, no matter where the spot was flashed in the receptive field (ON/OFF cells).

Two important aspects of retinal processing are already recognizable in this early work: lateral inhibition in space and differentiation in time. Because of the antagonistic action of the center and surround regions of the receptive field, ganglion cells respond strongly to stimuli whose intensity varies in space over the receptive field, such that center and surround are illuminated differently. And because the response to a light step lasts only a short time—typically tens of milliseconds to seconds—many ganglion cells seem to emphasize stimuli that change in time over static ones.

The visual world, of course, does not consist of spots and annuli. Thus, one needs to cast the stimulus–response relationship in a quantitative form that generalizes to arbitrary patterns of visual input. In its most general form, the stimulus is given by the intensity distribution $I(x, t, \lambda)$ on the retina, as a function of position x , time t , and wavelength λ . Under the above assumptions, the response of the retina consists of the firing rates $R_i(t)$ of each of its ganglion cells. To capture retinal processing, one thus seeks a mathematical function whose input is the stimulus $I(x, t, \lambda)$ and whose output is the time course of a ganglion cell's firing rate $R(t)$. This function will have a number of free parameters, which are optimized based on the measured responses to experimental stimuli. Finally, one can test the performance of this model with other types of stimuli. The following sections will illustrate some examples of this powerful approach.

Spatiotemporal Integration. Rodieck (1965) made an early and influential attempt at a quantitative description of cat ganglion cell responses. As observed earlier, a small spot of light flashed briefly on the receptive field center of an ON cell produced a brief increase in firing followed by an undershoot and gradual recovery of the baseline firing rate. The shape of this time course was approximated as

$$A(t) = \delta(t) - h e^{-t/\tau} \quad (1)$$

where $\delta(t)$ denotes the delta function pulse of firing and h is the size of the subsequent undershoot, which decays with time constant τ . As in previous experiments, the amplitude of this response depended on the location of the spot: large and positive in the center, small and negative in the surround, and zero somewhere in between. This spatial profile of the response amplitude was formalized as a “difference of Gaussians”

$$B(x) = k_c \cdot \exp\left(-\frac{x^2}{2r_c^2}\right) - k_s \cdot \exp\left(-\frac{x^2}{2r_s^2}\right) \quad (2)$$

where k_c and k_s are the amplitude of the center and the surround Gaussians and r_c and r_s are their respective radii. Thus, the change in the firing rate produced by flashing a spot at time $t = 0$ and location x is $A(t) \cdot B(x)$. Now any given light intensity pattern, $I(x, t)$, such as a white bar moved across the retina, can be decomposed into many small flashed spots. Rodieck's model assumed that the effects of all these spots simply sum up. Thus, the firing rate, $R(t)$, produced by the visual stimulus becomes

$$R(t) = R_0 + \iint I(x, t') \cdot B(x) \cdot A(t - t') dx dt' \quad (3)$$

where R_0 is the cell's maintained firing rate without stimulation. This expression can also be viewed as a cascade of a few simple transformations of the stimulus (Figure 1A). First, the stimulus $I(x, t)$ is summed over all space, with the weighting function $B(x)$. Then the resulting signal is passed through a filter with impulse response $A(t)$. The result is added to the baseline firing rate R_0 , and negative values R_0 of the resulting firing rate $R(t)$ are truncated to zero.

The parameters in this model, namely R_0 , $A(t)$, and $B(x)$, were derived from the flashing spot measurements. Then the model was tested using very different stimuli, consisting of various shapes moving steadily across the cell's receptive field. As Figure 1B shows, there was a remarkable correspondence between the observed time course of the firing rate and the predictions of the model.

This model of the light response is very attractive in its simplicity. For example, the time course of the response to a flash is identical no matter where in the receptive field the flash is presented, except for a scaling factor. This is termed “space–time separability” (Wandell, 1995), because the weighting function in equation 3 separates into a term depending only on time multiplied by a term depending only on space. Subsequent work showed that space–time separability is not quite satisfied in ganglion cell responses: for example, the response to light falling in the surround is delayed relative to the response in the center, owing to the time required for lateral signal flow through horizontal or amacrine cells, and transmission across an additional synapse (Enroth-Cugell and Freeman, 1987; Sakai and Naka, 1995; Benardete and Kaplan, 1997a). This led to a simple “modified difference-of-Gaussians model” (Figure 1C), in which light is pooled separately within the center and the surround; the two resulting signals are passed through two different filters, then summed to generate the firing rate (Enroth-Cugell et al., 1983; Dawis et al., 1984). For some ganglion cell types, center and surround also have a different spectral sensitivity, because they are fed by a different mix of photoreceptors. In general, the wavelength dependence of the retinal response is governed by the spectral sensitivities of the rods and cones, an aspect that varies a great deal among species. We will not elaborate on this topic here, but refer the reader to recent reviews of color processing (Wandell, 1995; Lee, 1996).

Another fundamental feature of Rodieck's model is the linearity of its response. Twice the intensity fluctuation will produce twice the firing rate fluctuation; more generally, the response to the sum of two intensity patterns is the sum of their individual responses, barring truncation in the final step of spike generation. Subsequently, it was found that a linear relationship between stimulus and firing rate holds only for some retinal ganglion cells and only under restricted conditions: the modulations of the light intensity must be small compared to the mean, and the range of these modulations must not change very much over time (Enroth-Cugell and Robson, 1966; Victor, 1987; Benardete and Kaplan, 1997b). These are very narrow constraints, and it now appears that under stimulus conditions resembling those of our natural visual experience, a linear description of retinal ganglion cell responses is of rather limited use.

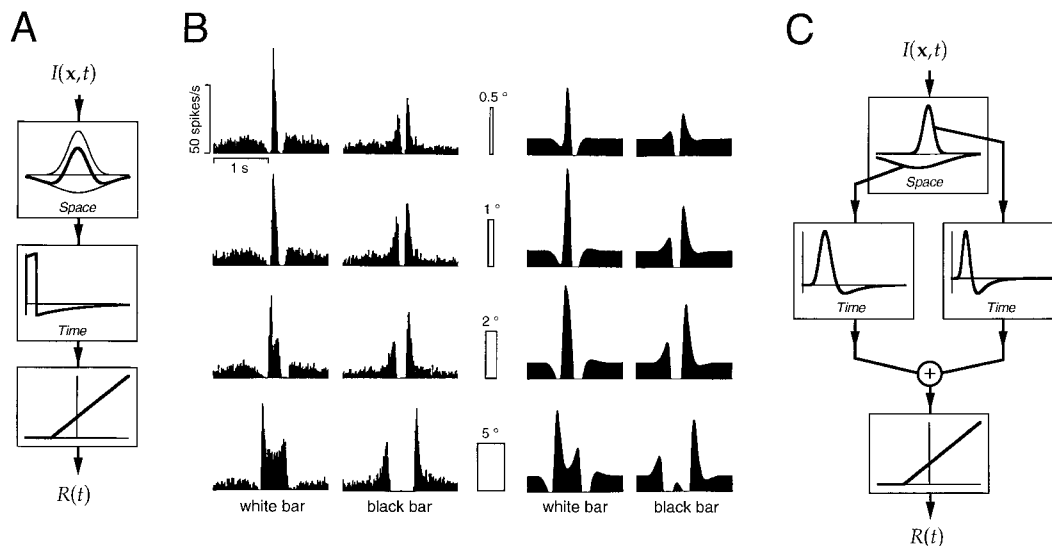


Figure 1. "Difference-of-Gaussians" Model to Predict the Firing Rate of a Cat Retinal Ganglion Cell

After Rodieck, 1965; Rodieck and Stone, 1965.

(A) At every time point, the intensity pattern on the retina is integrated over space with a weighting function that represents the receptive field profile (top, thick line). This profile is shaped as the difference of two concentric Gaussian surfaces (thin lines, see equation 2), here shown in a one-dimensional section through the center. The resulting signal is convolved in time with the retina's flash response (middle): a delta function (approximated in the graphic by a brief square pulse) followed by an exponential undershoot (equation 1). The result is added to a maintained firing rate and truncated (equation 3) to eliminate negative values (bottom).

The signal processing cascades described in this and subsequent figures contain three types of elements, and we will use the following graphic conventions to represent them: weighted spatial summation of light intensity is shown by the profile of the weighting function, with the horizontal axis labeled "Space" (top box in [A]); temporal filtering is represented by the impulse response of the filter, with the horizontal axis labeled "Time" (middle box in [A]); instantaneous transform of a signal is represented by a graph of output versus input, with the axes unlabeled (bottom box in [A]).

(B) Firing rate of an ON-type cat retinal ganglion cell in response to a bar swept across its receptive field (left; adapted from Rodieck and Stone, 1965), and as predicted by the model in (A) (right; adapted from Rodieck, 1965). The bars were either white or black on a gray background, varied in width from 0.5° to 5° (center), and were moved steadily at 10°/s.

(C) "Modified difference-of-Gaussians" model of the light response, in which center and surround are treated as separate pathways. Their Gaussian sensitivity curves may not be concentric, and their flash responses may differ. The resulting signals from both pathways are summed and rectified to produce the firing rate.

Nonlinear Processing. Victor (1987) has captured some of the nonlinear behavior of cat ganglion cells in a very successful model (Figure 2A): as in Rodieck's scheme, the light distribution is pooled linearly with a spatial weighting function and the result is passed through a temporal filter. However, the properties of this filter depend on its output. In particular, when the output is large—of either sign—the gain of the filter decreases and its waveform sharpens. With only a few parameters, this model accurately predicted the response to a variety of stimulus waveforms (Figure 2B), whereas any purely linear model produced large discrepancies. The net effect of this "contrast gain control" (Shapley and Victor, 1979, 1981) is that during large light fluctuations the retinal response is less sensitive and faster. This adjustment is very rapid: the time constant τ_c in Figure 2B was 15 ms, but a value of zero produced indistinguishable results.

This type of quantitative analysis has revealed that there are two very distinct types of ganglion cells in the cat retina. The so-called "X cells" (Figures 1 and 2) appear to integrate light from different points in space by simple weighted summation, while for the "Y cells" this is not the case (Enroth-Cugell and Robson, 1966; Hochstein and Shapley, 1976). The Y cell's receptive

field is several times larger than that of a nearby X cell. It is composed of many small spatial "subunits" that appear to process the stimulus independently (Victor and Shapley, 1979). Within the area of each subunit, the light intensity is integrated, again with an antagonistic receptive field and a biphasic impulse response (Figure 3). The result gets rectified, a highly nonlinear operation, and added to the output from all other subunits. This sum, after passing through another filter, specifies the firing rate (Victor, 1988). There may be as many as 100 such "nonlinear subunits" (Victor and Shapley, 1979). Due to their rectifying nature, a flashing spot anywhere within the receptive field can produce a burst of spikes at both onset and offset. Thus, the Y-type ganglion cell cannot signal the position of a small spot on the retina with the spatial resolution of an X cell. On the other hand, it is very sensitive to a fine textured pattern moving across the receptive field, since that induces intensity fluctuations for all the local subunits. The anatomical identity of the subunits is still uncertain. Their size is about equal to that of X cell centers, and they have been proposed to correspond to bipolar cells, with the rectification occurring in transmission to amacrine cells (Victor and Shapley, 1979).

The kind of processing performed by this model of

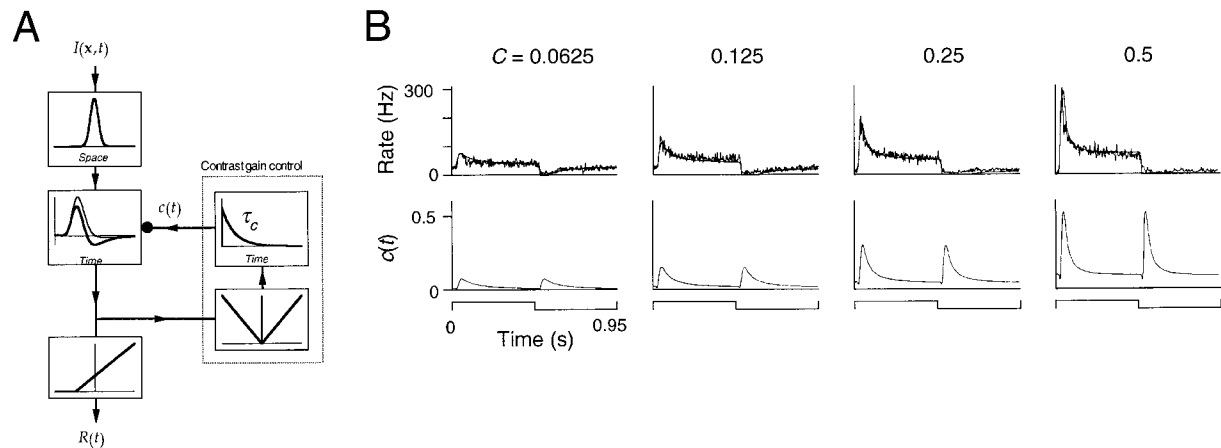


Figure 2. Cascade Model for Neural Coding by a Cat X Cell

After Victor, 1987.

(A) Only the pathway for the receptive field center is shown; see Figure 1A legend for conventions of this graphic shorthand. Light is integrated over the center's spatial profile. The result is passed through a band-pass temporal filter, then truncated to form the firing rate. To implement the contrast gain control, the output of the filter is full-wave rectified and averaged by a low-pass stage with time constant τ_c . The resulting signal, $c(t)$, is a neural measure of contrast and modifies the temporal processing properties: the flash response of the band-pass filter is more biphasic at high values of $c(t)$ (thick line) than at low values (thin line).

(B) Response of an ON-type X cell to contrast reversal (bottom trace) of a 1 cycle/degree sinusoidal grating at different modulation depths, C . Top plots show the measured firing rate (jagged line) and the prediction from the model in (A) (smooth line). Bottom plots show the neural measure of contrast, $c(t)$, that modulates the band-pass filter in (A). Note that the time course of the firing rate is more transient at large modulation depth, and this modification is reproduced accurately by the model.

the Y cell subunit—linear filtering followed by a nonlinear transformation and more linear filtering—is often called an LNL cascade (Hunter and Korenberg, 1986). This general scheme has been very useful in describing the response properties of other ganglion cells, for example, the P cells in macaque retina (Benardete and Kaplan, 1997b). Naka and colleagues have used the same formalism to model the behavior of ganglion cells in catfish retina (Korenberg et al., 1989). In a remarkable series of studies, they also recorded intracellularly from every major retinal cell type (Sakai and Naka, 1988). A quantitative assessment of their light responses revealed the contribution of each cell type to the ganglion cell response. In particular, the photoreceptors, horizontal cells, and bipolar cells produced essentially linear responses to light. Under the same stimulus conditions, amacrine cells showed strong nonlinear distortions, whose shape was distinct between the “sustained” and “transient” amacrine cells (Sakai and Naka, 1987). This supports the above speculation that retinal signals are strongly rectified during transmission to amacrine cells. It has also allowed the dissection of each ganglion cell's input into contributions from bipolars and the two types of amacrine cells (Sakai and Naka, 1995). Such experiments begin to give a biological identity to the cascade components in this formal description of the retinal code.

These examples illustrate that the relationship between light stimulus and firing rate can be phrased as a cascade model for many different types of retinal ganglion cells. This quantitative analysis has greatly illuminated the functional differences between classes of cells within the same retina. The X and Y cells, for example, were subsequently found to differ also in their morphological features and central projection patterns, reinforcing the impression that they play different roles in visual

processing (Stone, 1983; Wässle and Boycott, 1991). However, as discussed below, the present models are not yet satisfactory when faced with complex stimuli such as those encountered in natural vision. Furthermore, for some notable ganglion cell types, such as the direction-selective neurons found in many species (Barlow and Levick, 1965; Borg-Graham and Grzywacz, 1992), a quantitative model that predicts the response to a general ensemble of stimuli still remains to be found.

Precision

As is true for all other neurons, the response of retinal ganglion cells is stochastic, in that it varies somewhat even if one presents the identical visual stimulus on subsequent trials. Therefore, experimenters often repeat the same stimulus tens to hundreds of times and average the resulting responses to get an accurate estimate of the “true” firing rate. The brain, on the other hand, must interpret the ganglion cell output on a single trial. If the signal is carried by the neuron's firing rate, as is often assumed, then it can be decoded only to within the accuracy allowed by the trial-to-trial variation. On the other hand, the above models of the light response are all deterministic: they predict the average firing rate, as would be obtained from a large number of stimulus repeats, but make no statement about the range of the response one might see in a single trial. To understand what a ganglion cell can communicate, one must take into account not only the systematic “signal” but also the “noise” that corrupts it.

Many ganglion cells fire spontaneously even in darkness. In a sense, these spike trains represent pure noise, in absence of any stimulus-driven signal. Two extreme possibilities come to mind for the structure of such a spike train: it might be perfectly regular, with constant intervals between spikes as in pacemaker neurons; or




Figure 3. Neural Coding by the Cat Y Cell Subunits

After Hochstein and Shapley, 1976; Victor and Shapley, 1979; see Figure 1A legend for conventions of this graphic shorthand.

This pathway initiates in small subunits of the receptive field (thick profiles in top panel). Within each subunit's receptive field, light is integrated, the result is passed through a band-pass filter and then full-wave rectified. The rectified outputs from all subunits are pooled, passed through another linear filter, and converted to the firing rate. As for the X cell (Figure 2A), a contrast gain control modulates the temporal filter within each subunit (gray feedback pathway). The relevant contrast measure $c(t)$ is probably derived from the output of each subunit's rectifier and controls the contrast gain for both Y cells and X cells (Shapley and Victor, 1981). Note that the Y cell circuitry includes an additional signaling pathway, not elaborated in this figure, that produces a classical center and antagonistic surround (thin profiles in top panel). This pathway can be isolated by suitable visual stimuli (Victor and Shapley, 1979; Enroth-Cugell and Freeman, 1987) and appears to operate similarly to the cascade described for X cells (Figures 1C and 2A).

it might be perfectly random, with spikes produced at constant probability per unit time, regardless of the preceding firing history. This latter case corresponds to the so-called Poisson process and leads to an exponential distribution of inter-spike-intervals (Rieke et al., 1997). Retinal ganglion cells show intermediate statistics: their inter-spike-intervals are somewhat more regular than in a Poisson process, and their relative frequency is captured well by a gamma distribution (Kuffler et al., 1957; Levine and Shefner, 1977; Troy and Robson, 1992; Troy and Lee, 1994).

This inherent variability in the dark activity of retinal

ganglion cells must pose a limit to the processing of weak stimuli. Humans can detect very dim flashes of light that deliver only a handful of photons to the retina (Hecht et al., 1942; DeVries, 1943). In a landmark study, Barlow et al. (1971) asked how such weak stimuli are represented in the firing of retinal ganglion cells. They recorded action potentials from a ganglion cell in the cat eye, while stimulating the retina near the recording electrode with brief flashes of a small spot of light. In complete darkness, the ganglion cell was spontaneously active, firing about 20 spikes/s at irregular intervals (Figure 4A). Shortly after a flash, the rate of firing transiently increased and decreased again to the spontaneous rate. The total number of "extra" spikes fired in response to the flash was proportional to the flash intensity, with about 1–3 spikes generated for every photon absorbed by a retinal rod. How could one discriminate whether a flash occurred or not based on this ganglion cell's output? The authors conclude that one ought to simply count action potentials over a window of 200 ms (Figure 4A) and ask whether the count exceeded a given threshold. Due to the variability of the spontaneous activity, one obtains a wide distribution of spike counts in darkness; following a flash, the average number of spikes is larger, but again with a wide distribution (Figure 4B). For bright flashes, this distribution of spike counts has no overlap with the distribution in darkness, and one can set a threshold spike count such that flashes are detected perfectly. For dimmer flashes, the distributions do overlap and detection becomes imperfect. With a criterion of 50% correct detection and <2% false positives, the most sensitive ganglion cells reported flashes that delivered only 2–3 photons.

While the statistics of dark activity clearly affect the encoding of very dim lights, one might expect that stronger stimuli, which shape the pattern of firing substantially, also change the trial-to-trial variation of the response. In fact, it has been reported that the variance in the spike count over a given time window increases for stimuli that produce a larger mean spike count, with approximately a power law relationship (Levine et al., 1988, 1992). However, this conclusion is controversial, and other studies suggest that the response variation is essentially independent of the stimulus (van Dijk and Ringo, 1987; Troy and Robson, 1992; Croner et al., 1993). If so, then the noise in the response could simply be treated as a random value added to the spike count predicted from a deterministic model, such as the cascades discussed above.

It should be noted that these studies analyzed spike counts from ganglion cells over time windows ranging from 0.25 s to 1 s. The unstated assumption in this analysis is that whatever neuron interprets the ganglion cell output must integrate spikes for at least 0.25 s. This would effectively smear together retinal events that occurred in the preceding quarter of a second or more, a proposal that seems inconsistent with the speed of most visually guided behaviors. Primates clearly see with much higher time resolution: video clips presented at 4 frames/s appear as a sequence of snapshots, not a smoothly moving scene. Similarly cats, with their track record of catching birds in flight, must operate with better than 0.25 s time resolution. Thus, it is unclear

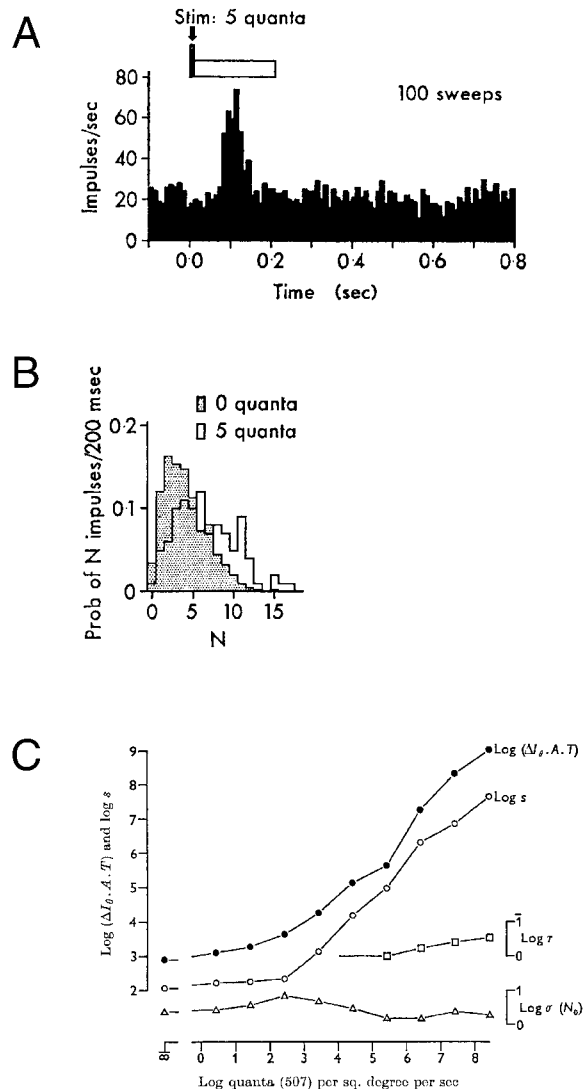


Figure 4. Encoding of Weak Flashes by a Ganglion Cell in the Cat Retina

After Barlow and Levick, 1969; Barlow et al., 1971.

(A) Firing rate before and after a flash that delivers an average of five photons to the cornea. The firing rate was averaged over 100 trials. The "extra" spikes caused by the flash correspond to the area of the peak above the maintained firing rate.

(B) To detect whether a flash occurred on an individual trial, spikes are counted over a 200 ms window (see box in [A]). This spike count shows considerable variation, but its probability distribution on trials with a flash is distinct from that on trials without a flash. An ideal detector of flashes would set a threshold somewhere between the modes of the two distributions and test whether the received spike count exceeds the threshold.

(C) Adaptation of this code to background light. Three factors can affect the detectability of flashes: the number of photons required to produce an extra spike (quantum-per-spike ratio s , open circles); the effective duration of the response (τ , open squares); and the standard deviation of the spike count (σ , open triangles). From these three parameters, one can compute the threshold intensity of a detectable flash (filled circles). All these quantities are plotted against the background light intensity, on axes logarithmic to base 10. Note that over a wide range the threshold intensity increases proportionally to the background and that this is mostly caused by a change in the quantum-per-spike ratio.

whether the variability in spike counts over long time windows relates in any way to the limits of visual processing. As discussed below, retinal ganglion cell spike trains are remarkably precise when analyzed on a finer time scale.

Adaptation

The neural code of retinal ganglion cells is not a static set of rules linking the visual stimulus to the optic nerve spike trains. Instead, it depends significantly on the overall properties of the visual scene. In particular, the code changes with the average light level by processes collectively termed "light adaptation" (Shapley and Enroth-Cugell, 1984).

The most pronounced effect of decreasing the mean light level is an increase in the sensitivity of retinal ganglion cells. For example, when the retina is stimulated by brief flashes on a steady background, the flash intensity required to elicit a criterion firing rate is lower for a dim background than for a bright background (Figure 4C). Over a wide range of intensities spanning several orders of magnitude, this threshold flash intensity is proportional to the mean light level (Barlow and Levick, 1969; Enroth-Cugell and Shapley, 1973a; Donner et al., 1990; Troy and Enroth-Cugell, 1993; Sakai and Naka, 1995). This so-called "Weber-Fechner law" of adaptation implies that the retina produces approximately the same response for two visual displays that are related by a simple proportional scaling of all intensity values.

Such behavior is of clear practical utility: because the intensity of the light illuminating the natural world changes over many orders of magnitude every day, so does the absolute intensity reflected by objects in the scene. However, the surface reflectance of these objects remains the same, and thus the relative ratios of intensities received from different parts of the scene are approximately independent of the illuminant. Through the process of light adaptation, the retina encodes the invariant features of objects and discards, for the most part, information about the absolute light level. Only a few ganglion cells appear dedicated to signaling the absolute intensity, for example to drive the reflex that keeps the pupil constricted in bright light (Barlow and Levick, 1969).

Along with the sensitivity, other aspects of the ganglion cell response change as well. In dim light, the time course of the response slows down considerably: a brief flash produces a burst of spikes with longer latency and longer duration (Enroth-Cugell and Shapley, 1973a; Donner et al., 1995). Thus, the ganglion cells integrate the visual input over a longer time interval before reporting it to the brain (Naka et al., 1979). This averaging may be required to attenuate the effects of neural noise under conditions where the signal is small, but it comes at the cost of impaired time resolution. Spatial integration by ganglion cells is also altered in dim light: the receptive field loses its antagonistic surround region (Barlow et al., 1957; Donner and Reuter, 1965; Stell et al., 1975; Masland and Ames, 1976; Bowling, 1980; Muller and Dacheux, 1997). As a result, the area in which light excites an ON-type ganglion cell expands somewhat. Again, this may be a strategy to enhance the sensory signal by collecting as much light as possible, at the expense of some spatial resolution (Atick and

Redlich, 1992). All these effects are also observed in human psychophysics (Shapley and Enroth-Cugell, 1984), and thus it appears that retinal processing largely accounts for the perceptual effects of light adaptation.

This adjustment of retinal sensitivity is remarkably fast, occurring within a few tenths of a second after the change in background intensity (Enroth-Cugell and Shapley, 1973a; Baylor and Hodgkin, 1974). In addition, a slower component requires several seconds to complete (Adelson, 1982; Nakatani et al., 1991; Yeh et al., 1996). At least three mechanisms contribute to the various aspects of light adaptation. First, individual photoreceptors adapt: their flash response becomes more sensitive and slower in dim light (Baylor and Hodgkin, 1974; Kraft et al., 1993; Yau, 1994). Most rods and cones show a range where sensitivity is inversely related to mean intensity, close to the ideal Weber law adaptation (Nakatani et al., 1991). At background intensities beyond this Weber regime, the receptor's flash response decreases much more rapidly, because of saturation in the photo-transduction process. At background intensities much below the Weber regime, the flash response is of constant amplitude. The underlying mechanisms of photoreceptor adaptation are becoming understood in molecular detail (Koutalos and Yau, 1996). Second, in dim light the retinal circuitry switches from cones to rods as the primary input neurons. Because rods are much more sensitive than cones, this contributes to the increased response sensitivity in dim light. In many species, this also produces a distinct shift in the spectral sensitivity of retinal ganglion cells, because rods and cones contain different photopigments. Finally, circuits postsynaptic to the receptor cells also alter their sensitivity during light adaptation. This is evident because ganglion cells already show light adaptation at intensities too low to trigger changes in individual rods (Donner et al., 1990). Furthermore, light falling on one portion of the receptive field center alters the cell's sensitivity to light in another portion (Schellart and Spekreijse, 1972; Enroth-Cugell and Shapley, 1973b). The mechanisms that give rise to this "network adaptation" still remain somewhat obscure.

On a formal level, all these effects of light adaptation can be viewed as a modulation of the stimulus-response relationship controlled by the mean light level. Several models of the ganglion cell light response can account for these effects quantitatively, by modulating the transfer properties of various elements in the cascade (Shapley and Enroth-Cugell, 1984; Purpura et al., 1990).

Summary

At this stage, it helps to summarize the commonly accepted elements of retinal coding. (1) The stimulus consists of the intensity distribution $I(x, t, \lambda)$ on the retina, as a function of position x , time t , and wavelength λ . The response of the retina consists of the firing rates $R_i(t)$ of each of its ganglion cells. (2) For each distinct type of ganglion cell, the expected firing rate $R(t)$ can be written as a functional of the stimulus $I(x, t, \lambda)$. Often, this relationship is expressed in the form of a cascade of linear filters and nonlinear transforms. Different cells of the same type process different parts of the scene with essentially the same stimulus-response relationship, although some parameters of this function, such as the receptive field size, may vary systematically with

retinal eccentricity. (3) On repeated presentations of the same stimulus, the actual measured firing rate varies stochastically about the expected firing rate $R(t)$. This random component of the firing rate is often thought to be independent of the stimulus, but it is unclear what retinal mechanisms determine its properties. (4) Light adaptation alters the function that relates stimuli to responses, depending on the mean level of illumination. Most notably this affects the gain of ganglion cells, but also their spatial receptive field, temporal integration, and spectral sensitivity.

New Aspects of Retinal Signaling

Recent work has suggested several revisions to this picture of the retinal code. They regard the relevant response variables, the nature of the relationship that links them to the stimulus, and the degree of plasticity in this relationship.

Ganglion Cells Are Not Independent Channels of Information

A central assumption in the work reviewed above is that the retinal code can be formulated by describing the responses of individual ganglion cells. An argument to support this assumption is that the dendritic fields of any given class of ganglion cells appear to "tile" the retina with little overlap, so that each point on the retina is covered by just one cell of that type (Wässle and Boycott, 1991). However, the ganglion cell's receptive field center tends to be somewhat larger than its dendritic field. For example, in cat retina, up to ten receptive field centers of ON-type X cells (and similarly for the OFF type) may overlap any given point (Peichl and Wässle, 1979). Thus, several neurons are conveying information about the same location in the visual image, and one needs to consider how their respective messages should be combined. Do they act as independent sources of information, such that each ganglion cell views the world through its receptive field window and makes an independent decision whether or not to fire? Or do these neurons somehow signal in a concerted fashion?

Evidence from multineuron recordings supports the latter view: retinal ganglion cells engage in significant patterns of concerted activity that cannot be derived from any single-neuron description. In particular, nearby cells of similar functional type have a strong tendency to fire in synchrony, much more frequently than expected by chance. This has been observed in many species, including tiger salamander, goldfish, rabbit, cat, and macaque (Arnett, 1978; Arnett and Spraker, 1981; Johnsen and Levine, 1983; Mastronarde, 1989; Meister, 1996; DeVries, 1999).

One can distinguish three types of concerted firing (Mastronarde, 1989) that differ in the time scale on which the spikes of two ganglion cells are synchronized. Recent work has identified the mechanisms by which they come about (Brivanlou et al., 1998). "Narrow" correlations involve synchrony of two ganglion cell spikes within 1 ms or less; these are caused by direct excitation among ganglion cells via electrical gap junctions. "Medium" correlations synchronize ganglion cells on a time scale of 10 ms (10–50 ms in salamander, 2–10 ms in cat); these arise from shared excitatory input the ganglion cells

receive from a common presynaptic neuron, likely an amacrine cell, again via gap junctions. "Broad" correlations produce synchrony on a scale of 50 ms (40–100 ms in salamander, 40–50 ms in cat) and result from a shared input signal that arrives from the outer retina through chemical synapses.

Among these, the "medium" correlations are most prevalent. In cat retina, these shared inputs cause about 80% of the maintained firing of Y cells (Mastrorade, 1983) and 15% for X cells. In salamander retina, they account for almost all the spontaneous activity of certain ganglion cells and for about 50% of all spikes recorded from a large population with an electrode array (Meister et al., 1995). These multineuron recordings also showed that firing is not restricted to pairs of ganglion cells but can involve groups of seven or more cells located within about 400 μm of each other (Schnitzer and Meister, 1996).

These correlations persist under many types of visual stimulation (Meister et al., 1995). Thus, one can subject each multicell pattern of synchronous spikes to the same response analysis previously applied to single neurons. The receptive field of a synchronous pair of spikes from two ganglion cells is generally smaller than the receptive fields of the two "parent" neurons and located at their intersection (Meister et al., 1995; Meister, 1996; Schnitzer and Meister, 1996). This is consistent with the above proposal that synchronous spikes arise in a shared presynaptic amacrine cell, which contributes part of the receptive field overlap of the two ganglion cells. It also implies that the synchronous spike pair (or, more generally, a multicell spike pattern) effectively signals the activity of the shared amacrine cell. Thus, the detection and proper interpretation of synchronous spikes by visual circuits in the brain could resolve the visual scene with a spatial grain finer than the ganglion cell population, for example at the resolution of small amacrine cell receptive fields.

Such a multiplexing scheme might serve to compress visual information for transmission through the optic nerve, where it is represented by the smallest number of neurons (Meister, 1996). At the first synapse in mammalian visual cortex, this number expands again by a factor of about 40 (Chow et al., 1950; Winfield and Powell, 1983); one suspects that the neural code for vision changes dramatically at this stage. These cells might decode the synchronous firing patterns among their afferent fibers (Alonso et al., 1996), if they are poised to fire only when several postsynaptic potentials superpose in time. While these proposals remain to be tested, synchronous firing is an important aspect of the retinal code simply because it affects much of the retinal output. Instead of single cell action potentials, one should consider the various multicell firing patterns as elementary symbols in retinal spike trains. Thus, the vocabulary by which the retina relates the visual scene may well be richer than estimated from single cell analysis.

Spike Trains from Retinal Ganglion Cells Can Be Very Precise and Reliable

As discussed above, it is often assumed that the brain must count many action potentials in order to reliably estimate a ganglion cell's firing rate and thus extract its visual message. However, a consideration of response

times in visually guided behavior suggests that during realistic vision this integration time should be no longer than about 50 ms. In this interval, the individual ganglion cell typically fires only a handful of spikes. Furthermore, there are situations where behavioral decisions are based on only a few ganglion cells. For example, our two-point acuity—the ability to distinguish two nearby spots from one—can resolve distances on the order of the spacing between cones in the fovea; under these conditions, the visual signal about the two spots is carried by just two ganglion cells. Furthermore, while we ponder this display, miniature eye movements continuously scan our retina over the image, so that each ganglion cell is activated for only about 20 ms at a time (Skavenski et al., 1979). Therefore, it appears that neural circuits in the brain must operate on ganglion cell responses containing only a few spikes or even single action potentials.

It is well known that strong visual stimuli can lead to very precise timing of individual action potentials (Levick and Zacks, 1970). For example, a step of light in the receptive field center of a cat retinal ganglion cell can reproducibly elicit spikes with a timing jitter of only 1–2 ms (Bolz et al., 1982). Recently, it has become clear that a much broader class of visual stimuli have this effect (Lankheet et al., 1989; Berry et al., 1997; Reich et al., 1997).

Rapid intensity transients appear to be a key stimulus feature for triggering precisely timed spikes. A useful experimental stimulus is given by random flicker: here, the intensity of a uniform field is updated at periodic short intervals, by drawing from a Gaussian probability distribution. This stimulus contains power distributed over a broad range of temporal frequencies. In many ganglion cells, it elicits spikes or brief bursts of spikes that are tightly locked to certain features in the stimulus sequence (Berry et al., 1997): a subsequent repeat of the same stimulus produces an almost identical spike train (Figure 5A). If one analyzes these responses by computing the spike time histogram averaged over many trials, one finds that the ganglion cell is perfectly silent most of the time, except for very brief periods (Figure 5A, bottom). At the onset of such an event, the firing rate rises to the maximal value within a few milliseconds and then rapidly drops to zero again.

Clearly, this behavior does not conform well to the conventional idea of the ganglion cell's coding variable: a maintained firing rate that is gradually modulated by the visual stimulus. Instead, the response consists of discrete firing events, each of which is specified sufficiently by the time of its first spike and the total number of spikes in the event (Figure 5B; Berry et al., 1997). These response variables are highly reliable: the trial-to-trial timing jitter of the first spike in an event is typically only 4 ms and sometimes less than 1 ms. The number of spikes in the event often varies by at most one from trial to trial (Berry et al., 1997; Berry and Meister, 1998). Thus, one concludes that individual ganglion cell action potentials are highly significant. Firing events differing by just one spike can reliably represent different visual features, and the timing of spikes reports the time of occurrence of such features with an accuracy of a few milliseconds.

What is the nature of the stimulus features that cause

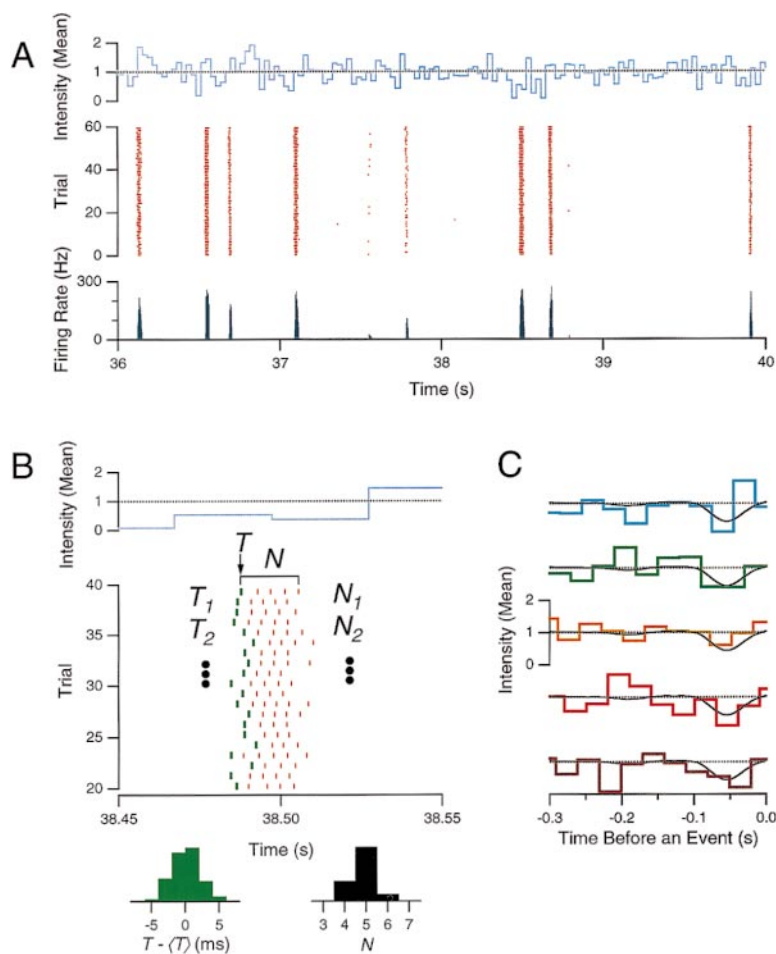


Figure 5. Precise Firing Events in the Response of a Salamander Retinal Ganglion Cell to Random Flicker

(A) The intensity of the uniformly illuminated field (top), a raster display of spikes on 60 repeated trials of this stimulus (middle), and the resulting firing rate derived from the peristimulus time histogram (bottom). The stimulus was generated by choosing the light intensity randomly from a Gaussian distribution every 30 ms. The mean intensity was in the photopic regime, and the standard deviation was 35% of the mean (for further experimental detail see Berry et al., 1997).

(B) A detail view of one firing event from (A). On each trial, the firing event is characterized by the time T of the first spike (thick marker in raster display) and the total number N of spikes fired. Bottom histograms show the distribution across trials for both quantities; $\langle T \rangle$ is the first-spike time averaged across trials.

(C) A sample of five intensity waveforms (thick lines) from the experiment in (A) that each produced a firing event with an average spike count $\langle N \rangle$ between 1.8 and 2.2 spikes. The average intensity waveform preceding all spikes from this cell is shown by the thin line (Berry and Meister, 1998).

firing? Many different episodes in a random flicker stimulus can reliably trigger the same kind of firing event (Figure 5C). It remains to be determined what property all these episodes have in common: this property could be identified as the visual message conveyed by the firing event. The collection of such messages from firing events in the ganglion cell population would constitute "what the eye tells the brain." Again, we should test our understanding of this code with explicit models that predict when such firing events occur for an arbitrary visual stimulus (Keat and Meister, 1997).

Neural responses with high temporal precision are not limited to strong stimulation. Even at very low contrasts—barely perceptible to the experimenter—the response consists of brief firing events with no maintained firing (Figure 6A). The precision of timing and spike number in these events declines only weakly with contrast (Berry et al., 1997). At zero contrast, namely steady illumination, the cells do exhibit stochastic firing, but the properties of this maintained activity bear little relevance to the response to time-varying stimuli. Precise firing also persists when the stimulus is spatially modulated, as in a randomly flickering checkerboard. In fact, ongoing experiments suggest that the retina operates in this same regime during natural vision. Video recordings of the natural environment projected onto the retina produced sparse and precise firing events in the ganglion cells of rabbit and salamander (M. J. B., J. Keat, and

M. M., unpublished data). In cells of the cat's lateral geniculate nucleus, visual stimuli with an intensity time course derived from natural scenes also elicited very precise firing events (Reinagel and Reid, 1998). Thus, it is possible that ganglion cell responses during natural vision are qualitatively different from those observed with traditional experimental stimuli such as drifting sinusoidal gratings.

Retinal Processing Adapts to Higher Stimulus Statistics

In order to encode the visual scene under many conditions, the retina must efficiently match the dynamic range of its output neurons to the range of sensory inputs. As discussed above, light adaptation accomplishes part of this by adjusting the retina's sensitivity to the prevailing intensity of light. However, the mean intensity is only one statistic that changes significantly during visual processing. For example, the range of light intensities, or "contrast," depends strongly on whether the scene is illuminated by direct light, casting sharp shadows, or illuminated diffusely by indirect light. If the retina maintained the same response–intensity relationship following an increase in contrast, its response would saturate over much of the range of inputs encountered in the new environment. Clearly, it would be beneficial to alter the response properties of ganglion cells, in order to adapt to the new range of intensities.

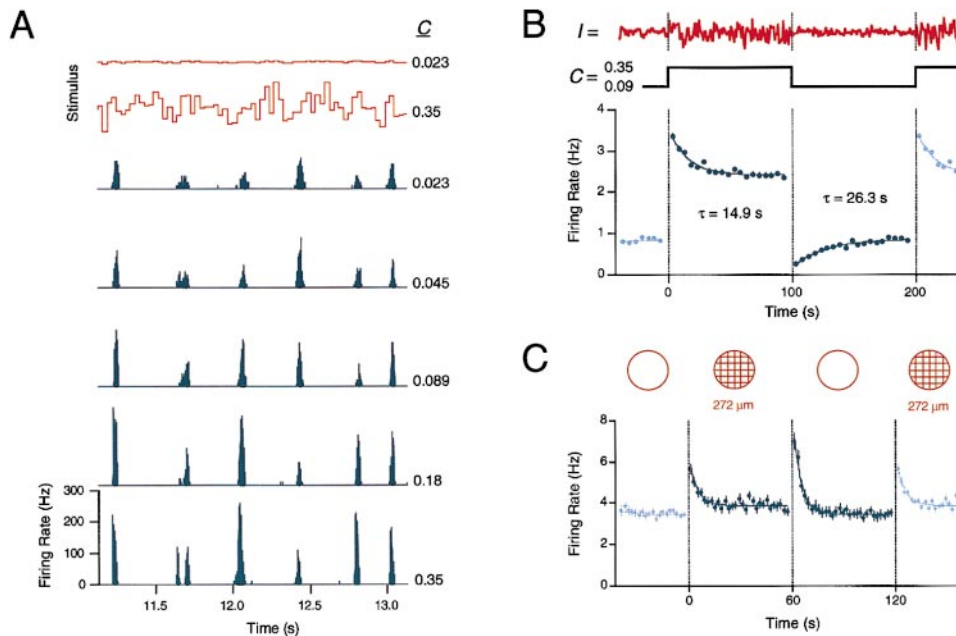


Figure 6. Contrast Adaptation in the Light Response of Salamander Retinal Ganglion Cells

(A) Firing events in response to uniform random flicker (see Figure 5A) of the same mean intensity but varying contrast. The contrast, C , measured as the standard deviation of the Gaussian intensity distribution in units of the mean, is indicated on the right. The two stimulus traces (top) show the time course of the light intensity for the two extreme values of contrast, $C = 0.023$ and $C = 0.35$. Note that the waveforms are identical except for the magnitude of fluctuations about the mean intensity.

(B) The average firing rate of a ganglion cell following a change in the contrast of uniform random flicker (from Smirnakis et al., 1997). Every 100 s, the contrast C alternated between 0.09 and 0.35. A sample intensity waveform is shown at the top, but each stimulus trial used a different random sequence. The firing rate was computed from the average spike count over 100 trials in 5 s time bins.

(C) The average firing rate of a ganglion cell following a change in the spatial pattern of random flicker (from Smirnakis et al., 1997). Every 60 s, the stimulus pattern alternated between a uniform field and a checkerboard of 0.272 mm square size. The uniform field and each square of the checkerboard were modulated by independent random flicker sequences, all with the same mean and a contrast C of 0.24. The firing rate was computed from the average spike count over 50 trials in 2 s time bins. Note that following adaptation to the checkerboard, the response is more sensitive to the uniform field, and vice versa, following adaptation to the uniform field, the response is more sensitive to the checkerboard.

Such adaptation to contrast is a well-known phenomenon in human psychophysics. For example, after prolonged viewing of a high-contrast grating, our sensitivity for detection of similar gratings is much reduced (Blakemore and Campbell, 1969). The loss of sensitivity under high-contrast conditions and subsequent recovery under low-contrast conditions requires several seconds to tens of seconds (Blakemore and Campbell, 1969; Schieting and Spillmann, 1987). Similar effects of slow contrast adaptation have been observed in the responses of neurons in the visual cortex (Albrecht et al., 1984; Ohzawa et al., 1985; Allison et al., 1993; Carandini and Ferster, 1997). These changes have been thought to arise in cortical circuitry, because adapting stimuli presented to one eye can have some effect on test stimuli presented to the other eye. However, recent work shows that such gradual contrast adaptation also occurs in the retina and substantially modifies the light response of retinal ganglion cells.

These changes are revealed in experiments where a flickering visual stimulus suddenly changes from low contrast to high contrast of the same mean intensity (Donner et al., 1991; Smirnakis et al., 1997). Following such a switch, the retina gradually reduces its sensitivity. When this adaptation is complete, one finds that in the

high-contrast environment larger intensity transients are required to trigger a ganglion cell response than in the low-contrast environment. For some cells, the threshold for firing appears to vary in proportion to the contrast (Sakai et al., 1995); such neurons detect and report intensity fluctuations normalized to the average range of fluctuations (Figure 6A). After a shift back to low contrast, the retina's sensitivity gradually recovers to the initial state.

These adjustments in sensitivity are relatively slow and develop with exponential time constants on the order of 5–20 s (Figure 6B). Interestingly, the decay of sensitivity after the switch to a high-contrast environment occurs several times more rapidly than the recovery in a low-contrast environment (Smirnakis et al., 1997). This long time course of contrast adaptation distinguishes the process from the “contrast gain control” discussed above, which relies on a nearly instantaneous negative feedback effect during strong light transients (Shapley and Victor, 1978; Victor, 1987). Instead, the time course of contrast adaptation in retinal ganglion cells resembles that observed psychophysically and represents a slow adjustment of the retinal code to a change in the statistical makeup of the environment. It has been argued that a reliable assessment of second

order statistics, such as the range of intensity fluctuations about the mean, requires many independent samples of the light intensity, and thus adaptation to the width of the intensity distribution is necessarily slower than to the mean (Donner et al., 1990). Furthermore, a sudden increase in the intensity range rapidly produces very large or very small intensity values and, thus, should be detected more readily than a sudden decrease, which might explain the observed asymmetry in adaptation to high and low contrast (DeWeese and Zador, 1998).

The neurons that sense contrast and control this adaptation appear to integrate intensity fluctuations over a spatial area of several hundred micrometers on the retina. In fact, there are at least two different contrast sensors with different spatial sensitivities (Smirnakis et al., 1997). This allows retinal ganglion cells to adapt to an environment of low spatial frequency while remaining very sensitive to stimuli of high spatial frequency and vice versa (Figure 6C). From these observations, one can conclude that contrast adaptation is not implemented by the individual photoreceptor (Donner et al., 1991; Smirnakis et al., 1997). Intracellular recordings from retinal neurons in catfish have shown directly that the response properties of photoreceptors, horizontal cells, and bipolar cells do not change when the contrast of a flicker stimulus is altered (Sakai and Naka, 1987). On the other hand, visual processing by amacrine and ganglion cells is altered substantially. Thus, the site of contrast adaptation may lie in the inner retina, but its cellular mechanisms remain unknown.

What might be the functional benefit of this process? As suggested above, it may represent a dynamic adjustment of the retina's output range to the range of its input signals. This would avoid saturation of the ganglion cell response under high-contrast conditions while still allowing a high signal-to-noise ratio under low-contrast conditions. On the other hand, contrast adaptation also plays a "computational" role: it generates a neural representation of the visual scene that has been normalized to the average size of intensity fluctuations in the scene. In the process, information about the absolute contrast in the image is discarded (Figure 6A). In an analogous fashion, light adaptation normalizes the neural image to the mean light level and discards information about the absolute intensity in the image. The purpose of both operations would be to remove from the image behaviorally uninteresting aspects that are mostly dependent on the conditions of illumination or the average structure of the environment, while preserving and emphasizing the differences between objects in the visual scene. One can speculate further that each successive stage of the early visual system adapts to—and consequently discards—what appear to be constants in the neural representation from the previous stage (Barlow, 1990).

Future Directions

Recent work on the retina has documented several aspects of visual processing, such as concerted coding and contrast adaptation, that previously were thought to occur only at the level of the cortex. This emphasizes that the retinal transformation of the visual stimulus is more sophisticated than the simple linear filtering one

still finds portrayed in textbooks. One important goal is to better understand the cellular mechanisms underlying such phenomena.

Natural Vision

However, even at the schematic level of the retinal models discussed above, we do not yet have a satisfying understanding of how the retina functions during vision. To date we cannot predict—to any acceptable degree of accuracy—the firing of ganglion cells while an animal is viewing a natural scene. Similarly, no one would venture to describe the visual scene given only a recording of the optic nerve's spike trains. The reason for this state of affairs lies in the fact that natural vision has not been the focus of much retinal research. The late 1960's saw the introduction to visual neurophysiology of powerful methods from systems analysis: the use of sinusoid visual stimuli, a systematic variation of the stimulus parameters, and an emphasis on testing retinal processing with weak responses just above threshold (Stone, 1983; Shapley and Lam, 1993). These approaches have been very successful in defining different retinal pathways and identifying their temporal and spatial receptive fields and transfer functions. However, the resulting models of retinal processing do not generalize well to real vision.

Natural vision happens almost exclusively far above the response threshold. Outside of psychophysics dark-rooms, we are rarely placed in a situation where something important is barely detectable. Real visual scenes contain intensities spanning several orders of magnitude (van der Schaaf and van Hateren, 1996), with variation on all spatial and temporal scales (Barlow, 1961; Field, 1987; Dong and Atick, 1995). This is likely to influence the neural code significantly: for example, as discussed above, the broad temporal spectrum can lead to precisely timed firing events in retinal ganglion cells, which are poorly captured by the prevalent models for a continuous ganglion cell firing rate. In summary, while questions phrased about vision near the detection threshold may be more tractable, it is not clear that the answers are relevant to everyday vision. As a result, we have an exquisite understanding of how a toad can perceive meal worms in almost complete darkness (Aho et al., 1993a, 1993b), but only the most qualitative notion of what our own retina does while we read this article.

The only effective remedy will be to study visual processing under conditions of natural stimulation. There has been a general reluctance to use natural images or movies in vision research, mostly due to the seemingly intractable complexity of natural scenes, the need to consider the animal's eye movements, and the obvious bias that results from choosing any one stimulus from such a large set. On the other hand, given the large uncertainties about what actually happens during natural vision, studying the response to even one or a few brief epochs of real visual input will likely be revealing. A number of recent developments have made the associated problems more tractable: current video technology greatly facilitates the acquisition, editing, and display of natural stimuli (Land, 1992; Gallant et al., 1998). Although standard video cannot perfectly reproduce the resolution and intensity range of natural scenes, it should be remembered that television in fact provides much of the daily visual input for human primates. Also,

there has been considerable progress in understanding the statistics of natural scenes. Images from the natural environment are not random, but contain very strong regularities in the spatial, temporal, and chromatic domains (Field, 1987; Ruderman, 1994; Dong and Atick, 1995; van der Schaaf and van Hateren, 1996; Ruderman et al., 1998). Such insights about the composition of natural scenes will also help in developing synthetic stimuli whose image statistics at least approximate the natural world (Field, 1994; Ruderman, 1997).

Information Theory

Finally, the analysis of retinal processing will benefit greatly from the fruitful application of information theory. It has long been recognized that the essential substance transmitted by neurons is not electric charge or neurochemicals, but information (MacKay and McCulloch, 1952). In analyzing a neural system, it is essential to measure and track the flow of this substance, just as in studies of the vascular system one might want to measure blood flow. Fortunately, there exists a well established formalism by which to measure information and analyze its processing (Shannon and Weaver, 1963; Rieke et al., 1997). A recent study implemented a specific model of retinal coding to attempt a reconstruction of the visual stimulus directly from the spike trains of salamander ganglion cells (Warland et al., 1997). A local group of four cells from different functional types was found to transmit information about the visual stimulus at approximately 8 bits/s. This can be interpreted as resolving $2^8 = 256$ different visual features every second and, thus, provides at least one answer to the question posed at the outset of this review. It is also possible to measure the information content of spike trains in a model-independent fashion, which requires no knowledge or assumptions about the stimulus-response relationship (Strong et al., 1998). Remarkably, this latter method yields information rates about three times higher (M. J. B. and M. M., unpublished data), suggesting that the current model captures only about one third of what the ganglion cells really communicate. Such an objective measure is essential in assessing how close we come to understanding what the neural code is.

Information theory also offers an understanding of why the neural code is the way it is. At a simple level, one may regard the retina as a device that attempts to compress the visual scenes provided by the natural world so they can be transmitted through a small number of fibers of limited dynamic range. To this end, it would be important to avoid redundancy, namely a situation where several optic nerve fibers are reliably communicating the same signal. Natural images contain a great deal of redundancy, for example, because nearby points of the same object tend to have the same intensity. It could be beneficial to process the image with a simple spatial filter before encoding it (Attneave, 1954; Barlow, 1961). Given the actual image statistics of the natural world, one can formalize this problem and calculate the spatial filter that would optimally reduce the redundancy between different points in the image; by similar arguments, one can propose filters that reduce redundancy in the time domain and in the spectral domain. Remarkably, this optimal solution correctly predicts the dominant features of retinal ganglion cell receptive fields:

lateral inhibition in space, band-pass filtering in time, and even the observed spectral sensitivity curves (Srinivasan et al., 1982; Buchsbaum and Gottschalk, 1983; Atick, 1992; Atick et al., 1992; Dan et al., 1996). At low light levels, when visual processing is seriously affected by noise, a certain degree of redundancy is beneficial to safeguard reliable transmission. Again, one can compute how the optimal spatial filter should change depending on the signal-to-noise ratio, and the result predicts correctly the measured profiles of ganglion cell receptive fields at different light levels (Atick and Redlich, 1992; van Hateren, 1992). Given the simplicity of the assumptions, it is surprising how well this interpretation works. In fact, this is currently the only framework that predicts the functional properties of retinal ganglion cells from a simple underlying principle. It seems compelling to conclude that data compression is truly an important purpose of retinal circuitry and that evolution has adapted its strategies to the visual ensemble provided by the environment. One hopes that as we elaborate other aspects of retinal coding, they will find a similarly satisfying interpretation.

References

- Adelson, E.H. (1982). Saturation and adaptation in the rod system. *Vision Res.* 22, 1299–1312.
- Aho, A.C., Donner, K., and Reuter, T. (1993a). Retinal origins of the temperature effect on absolute visual sensitivity in frogs. *J. Physiol.* 463, 501–521.
- Aho, A.C., Donner, K., Helenius, S., Larsen, L.O., and Reuter, T. (1993b). Visual performance of the toad (*Bufo bufo*) at low light levels: retinal ganglion cell responses and prey-catching accuracy. *J. Comp. Physiol. [A]* 172, 671–682.
- Albrecht, D.G., Farrar, S.B., and Hamilton, D.B. (1984). Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. *J. Physiol.* 347, 713–739.
- Allison, J.D., Casagrande, V.A., Debruyn, E.J., and Bonds, A.B. (1993). Contrast adaptation in striate cortical neurons of the nocturnal primate bushy baby (*Galago crassicaudatus*). *Vis. Neurosci.* 10, 1129–1139.
- Alonso, J.M., Usrey, W.M., and Reid, R.C. (1996). Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature* 383, 815–819.
- Arnett, D.W. (1978). Statistical dependence between neighboring retinal ganglion cells in goldfish. *Exp. Brain Res.* 32, 49–53.
- Arnett, D., and Spraker, T.E. (1981). Cross-correlation analysis of the maintained discharge of rabbit retinal ganglion cells. *J. Physiol.* 317, 29–47.
- Atick, J.J. (1992). Could information theory provide an ecological theory of sensory processing? *Network* 3, 213–251.
- Atick, J.J., and Redlich, A.N. (1992). What does the retina know about natural scenes? *Neural Comput.* 4, 196–210.
- Atick, J.J., Li, Z., and Redlich, A.N. (1992). Understanding retinal color coding from first principles. *Neural Comput.* 4, 559–572.
- Attneave, F. (1954). Some informational aspects of visual perception. *Psychol. Rev.* 61, 183–193.
- Barlow, H.B. (1953). Summation and inhibition in the frog's retina. *J. Physiol.* 119, 69–88.
- Barlow, H.B. (1961). Possible principles underlying the transformation of sensory messages. In *Sensory Communication*, W.A. Rosenblith, ed. (Cambridge, MA: MIT Press), pp. 217–234.
- Barlow, H.B. (1990). A theory about the functional role and synaptic mechanism of visual after-effects. In *Vision: Coding and Efficiency*, C. Blakemore, ed. (Cambridge: University Press), pp. 363–375.
- Barlow, H.B., and Levick, W.R. (1965). The mechanism of directionally selective units in rabbit's retina. *J. Physiol.* 178, 477–504.

- Barlow, H.B., and Levick, W.R. (1969). Changes in the maintained discharge with adaptation level in the cat retina. *J. Physiol.* 202, 699–718.
- Barlow, H.B., and Levick, W.R. (1969). Three factors limiting the reliable detection of light by retinal ganglion cells of the cat. *J. Physiol.* 200, 1–24.
- Barlow, H.B., Fitzhugh, R., and Kuffler, S.W. (1957). Change of organization of the receptive fields of the cat's retina during dark adaptation. *J. Physiol.* 137, 338–354.
- Barlow, H.B., Levick, W.R., and Yoon, M. (1971). Responses to single quanta of light in retinal ganglion cells of the cat. *Vision Res.* 3, 87–101.
- Baylor, D.A., and Hodgkin, A.L. (1974). Changes in time scale and sensitivity in turtle photoreceptors. *J. Physiol.* 242, 729–758.
- Benardete, E.A., and Kaplan, E. (1997a). The receptive field of the primate P retinal ganglion cell, I: Linear dynamics. *Vis. Neurosci.* 14, 169–185.
- Benardete, E.A., and Kaplan, E. (1997b). The receptive field of the primate P retinal ganglion cell, II: nonlinear dynamics. *Vis. Neurosci.* 14, 187–205.
- Berry, M.J., and Meister, M. (1998). Refractoriness and neural precision. *J. Neurosci.* 18, 2200–2211.
- Berry, M.J., Warland, D.K., and Meister, M. (1997). The structure and precision of retinal spike trains. *Proc. Nat. Acad. Sci. USA* 94, 5411–5416.
- Blakemore, C., and Campbell, F.W. (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *J. Physiol.* 203, 237–260.
- Bolz, J., Rosner, G., and Wässle, H. (1982). Response latency of brisk-sustained (X) and brisk-transient (Y) cells in the cat retina. *J. Physiol.* 328, 171–190.
- Borg-Graham, L.J., and Grzywacz, N.M. (1992). A model of the directional selectivity circuit in retina: transformations by neurons singly and in concert. In *Single Neuron Computation*, T. McKenna, J. Davis, and S.F. Zornetzer, eds. (New York: Academic Press), pp. 347–375.
- Bowling, D.B. (1980). Light responses of ganglion cells in the retina of the turtle. *J. Physiol.* 299, 173–196.
- Brivanlou, I.H., Warland, D.K., and Meister, M. (1998). Mechanisms of concerted firing among retinal ganglion cells. *Neuron* 20, 527–539.
- Buchsbaum, G., and Gottschalk, A. (1983). Trichromacy, opponent colours coding and optimum colour information transmission in the retina. *Proc. R. Soc. Lond. B Biol. Sci.* 220, 89–113.
- Carandini, M., and Ferster, D. (1997). A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science* 276, 949–952.
- Chow, K.-L., Blum, J.S., and Blum, R.A. (1950). Cell ratios in the thalamo-cortical visual system of *Macaca mulatta*. *J. Comp. Neurol.* 92, 227–239.
- Croner, L.J., Purpura, K., and Kaplan, E. (1993). Response variability in retinal ganglion cells of primates. *Proc. Natl. Acad. Sci. USA* 90, 8128–8130.
- Dan, Y., Atick, J.J., and Reid, R.C. (1996). Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J. Neurosci.* 16, 3351–3362.
- Dawis, S., Shapley, R., Kaplan, E., and Tranchina, D. (1984). The receptive field organization of X-cells in the cat: spatiotemporal coupling and asymmetry. *Vision Res.* 24, 549–564.
- DeVries, H.L. (1943). The quantum character of light and its bearing upon threshold of vision, the differential sensitivity and visual acuity of the eye. *Physica* 10, 553–564.
- DeVries, S.H. (1999). Correlated firing in rabbit retinal ganglion cells. *J. Neurophysiol.*, in press.
- DeVries, S.H., and Baylor, D.A. (1997). Mosaic arrangement of ganglion cell receptive fields in rabbit retina. *J. Neurophysiol.* 78, 2048–2060.
- DeWeese, M., and Zador, A. (1998). Asymmetric dynamics in optimal variance adaptation. *Neural Comput.* 10, 1179–1202.
- Dong, D.W., and Atick, J.J. (1995). Statistics of natural time-varying images. *Network* 6, 345–358.
- Donner, K.O., and Reuter, T. (1965). The dark-adaptation of single units in the frog's retina and its relation to the regeneration of rhodopsin. *Vision Res.* 5, 615–632.
- Donner, K., Copenhagen, D.R., and Reuter, T. (1990). Weber and noise adaptation in the retina of the toad *Bufo marinus*. *J. Gen. Physiol.* 95, 733–753.
- Donner, K., Djupsund, K., Reuter, T., and Vaisanen, I. (1991). Adaptation to light fluctuations in the frog retina. *Neurosci. Res. Suppl.* 15, S175–S184.
- Donner, K., Koskelainen, A., Djupsund, K., and Hemila, S. (1995). Changes in retinal time scale under background light: observations on rods and ganglion cells in the frog retina. *Vision Res.* 35, 2255–2266.
- Dowling, J.E. (1987). *The Retina: An Approachable Part of the Brain* (Cambridge, MA: Harvard University Press).
- Enroth-Cugell, C., and Freeman, A.W. (1987). The receptive-field spatial structure of cat retinal Y cells. *J. Physiol.* 384, 49–79.
- Enroth-Cugell, C., and Robson, J.G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *J. Physiol.* 187, 517–552.
- Enroth-Cugell, C., and Shapley, R.M. (1973a). Adaptation and dynamics of cat retinal ganglion cells. *J. Physiol.* 233, 271–309.
- Enroth-Cugell, C., and Shapley, R.M. (1973b). Flux, not retinal illumination, is what cat retinal ganglion cells really care about. *J. Physiol.* 233, 311–326.
- Enroth-Cugell, C., Robson, J.G., Schweitzer-Tong, D.E., and Watson, A.B. (1983). Spatio-temporal interactions in cat retinal ganglion cells showing linear spatial summation. *J. Physiol.* 341, 279–307.
- Field, D.J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A* 4, 2379–2394.
- Field, D.J. (1994). What is the goal of sensory coding? *Neural Comput.* 6, 559–601.
- Gallant, J.L., Connor, C.E., and Van Essen, D.C. (1998). Neural activity in areas V1, V2 and V4 during free viewing of natural scenes compared to controlled viewing. *Neuroreport* 9, 2153–2158.
- Hecht, S., Schlaer, S., and Pirenne, M. (1942). Energy, quanta, and vision. *J. Gen. Physiol.* 25, 819–840.
- Hochstein, S., and Shapley, R.M. (1976). Linear and nonlinear spatial subunits in Y cat retinal ganglion cells. *J. Physiol.* 262, 265–284.
- Hunter, I.W., and Korenberg, M.J. (1986). The identification of nonlinear biological systems: Wiener and Hammerstein cascade models. *Biol. Cybern.* 55, 135–144.
- Johnsen, J.A., and Levine, M.W. (1983). Correlation of activity in neighbouring goldfish ganglion cells: relationship between latency and lag. *J. Physiol.* 345, 439–449.
- Keat, J.R., and Meister, M. (1997). Prediction of precisely timed spike events in retinal ganglion cells. *Soc. Neurosci. Abstr.* 23, 1960.
- Korenberg, M.J., Sakai, H.M., and Naka, K. (1989). Dissection of the neuron network in the catfish inner retina. III. Interpretation of spike kernels. *J. Neurophysiol.* 61, 1110–1120.
- Koutalos, Y., and Yau, K.W. (1996). Regulation of sensitivity in vertebrate rod photoreceptors by calcium. *Trends Neurosci.* 19, 73–81.
- Kraft, T.W., Schneeweis, D.M., and Schnapf, J.L. (1993). Visual transduction in human rod photoreceptors. *J. Physiol.* 464, 747–765.
- Kuffler, S.W. (1953). Discharge patterns and functional organization of mammalian retina. *J. Neurophysiol.* 16, 37–68.
- Kuffler, S.W., Fitzhugh, R., and Barlow, H.B. (1957). Maintained activity in the cat's retina in light and darkness. *J. Gen. Physiol.* 40, 683–702.
- Land, M.F. (1992). Predictable eye-head coordination during driving. *Nature* 359, 318–320.
- Lankheet, M.J., Molenaar, J., and van de Grind, W.A. (1989). The spike generating mechanism of cat retinal ganglion cells. *Vision Res.* 29, 505–517.
- Lee, B.B. (1996). Receptive field structure in the primate retina. *Vision Res.* 36, 631–644.
- Levick, W.R., and Zacks, J.L. (1970). Responses of cat retinal ganglion cells to brief flashes of light. *J. Physiol.* 206, 677–700.

- Levine, M.W., and Shefner, J.M. (1977). A model for the variability of interspike intervals during sustained firing of a retinal neuron. *Biophys. J.* 19, 241–252.
- Levine, M.W., Zimmerman, R.P., and Carrion-Carere, V. (1988). Variability in responses of retinal ganglion cells. *J. Opt. Soc. Am. A* 5, 593–597.
- Levine, M.W., Cleland, B.G., and Zimmerman, R.P. (1992). Variability of responses of cat retinal ganglion cells. *Vis. Neurosci.* 8, 277–279.
- MacKay, D.M., and McCulloch, W.S. (1952). The limiting information capacity of a neuronal link. *Bull. Math. Biophys.* 14, 127–135.
- Masland, R.H., and Ames, A. (1976). Responses to acetylcholine of ganglion cells in an isolated mammalian retina. *J. Neurophysiol.* 39, 1220–1235.
- Mastrorade, D.N. (1983). Correlated firing of cat retinal ganglion cells. I. Spontaneously active inputs to X- and Y-cells. *J. Neurophysiol.* 49, 303–324.
- Mastrorade, D.N. (1989). Correlated firing of retinal ganglion cells. *Trends Neurosci.* 12, 75–80.
- Meister, M. (1996). Multineuronal codes in retinal signaling. *Proc. Natl. Acad. Sci. USA* 93, 609–614.
- Meister, M., Lagnado, L., and Baylor, D.A. (1995). Concerted signaling by retinal ganglion cells. *Science* 270, 1207–1210.
- Muller, J.F., and Dacheux, R.F. (1997). Alpha ganglion cells of the rabbit retina lose antagonistic surround responses under dark adaptation. *Vis. Neurosci.* 14, 395–401.
- Naka, K.I., Chan, R.Y., and Yasui, S. (1979). Adaptation in catfish retina. *J. Neurophysiol.* 42, 441–454.
- Nakatani, K., Tamura, T., and Yau, K.W. (1991). Light adaptation in retinal rods of the rabbit and two other nonprimate mammals. *J. Gen. Physiol.* 97, 413–435.
- Ohzawa, I., Sclar, G., and Freeman, R.D. (1985). Contrast gain control in the cat's visual system. *J. Neurophysiol.* 54, 651–667.
- Peichl, L., and Wässle, H. (1979). Size, scatter and coverage of ganglion cell receptive field centres in the cat retina. *J. Physiol.* 297, 117–141.
- Perkel, D.H., and Bullock, T.H. (1968). Neural coding. *Neurosci. Res. Program Bull.* 6, 221–348.
- Purpura, K., Tranchina, D., Kaplan, E., and Shapley, R.M. (1990). Light adaptation in the primate retina: analysis of changes in gain and dynamics of monkey retinal ganglion cells. *Vis. Neurosci.* 4, 75–93.
- Reich, D.S., Victor, J.D., Knight, B.W., Ozaki, T., and Kaplan, E. (1997). Response variability and timing precision of neuronal spike trains in vivo. *J. Neurophysiol.* 77, 2836–2841.
- Reinagel, P., and Reid, R.C. (1998). Visual stimulus statistics and the reliability of spike timing in the LGN. *Soc. Neurosci. Abstr.* 24, 139.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R.R., and Bialek, W. (1997). *Spikes: Exploring the Neural Code* (Cambridge, MA: MIT Press).
- Rodieck, R.W. (1965). Quantitative analysis of cat retinal ganglion cell response to visual stimuli. *Vision Res.* 5, 583–601.
- Rodieck, R.W. (1998). *The first steps in seeing* (Sunderland, MA: Sinauer Associates).
- Rodieck, R.W., and Stone, J. (1965). Response of cat retinal ganglion cells to moving visual patterns. *J. Neurophysiol.* 28, 819–832.
- Rodieck, R.W., Brening, R.K., and Watanabe, M. (1993). The origin of parallel visual pathways. In *Contrast Sensitivity*, R. Shapley and D.M.-K. Lam, eds. (Cambridge, MA: MIT Press).
- Ruderman, D.L. (1994). The statistics of natural images. *Network* 5, 517–548.
- Ruderman, D.L. (1997). Origins of scaling in natural images. *Vision Res.* 37, 3385–3398.
- Ruderman, D.L., Cronin, T.W., and Chiao, C.-C. (1998). Statistics of cone responses to natural images: implications for visual coding. *J. Opt. Soc. Am. A* 15, 2036–2045.
- Sakai, H.M., and Naka, K. (1987). Signal transmission in the catfish retina. V. Sensitivity and circuit. *J. Neurophysiol.* 58, 1329–1350.
- Sakai, H.M., and Naka, K.-I. (1988). Neuron network in catfish retina. *Prog. Ret. Res.* 7, 149–208.
- Sakai, H.M., and Naka, K. (1995). Response dynamics and receptive-field organization of catfish ganglion cells. *J. Gen. Physiol.* 105, 795–814.
- Sakai, H.M., Wang, J.L., and Naka, K. (1995). Contrast gain control in the lower vertebrate retinas. *J. Gen. Physiol.* 105, 815–835.
- Schellart, N.A., and Spekreijse, H. (1972). Dynamic characteristics of retinal ganglion cell responses in goldfish. *J. Gen. Physiol.* 59, 1–21.
- Schieting, S., and Spillmann, L. (1987). Flicker adaptation in the peripheral retina. *Vision Res.* 27, 277–284.
- Schnitzer, M.J., and Meister, M. (1996). Multi-neuronal firing patterns among retinal ganglion cells carry distinct visual messages. *Soc. Neurosci. Abstr.* 22, 494.
- Shannon, C.E., and Weaver, W. (1963). *The Mathematical Theory of Communication*, Second Edition (Chicago: University of Illinois Press).
- Shapley, R., and Enroth-Cugell, C. (1984). Visual adaptation and retinal gain controls. *Prog. Ret. Res.* 3, 263–346.
- Shapley, R., and Lam, D.M.-K. (1993). *Contrast Sensitivity* (Cambridge, MA: MIT Press).
- Shapley, R.M., and Victor, J.D. (1978). The effect of contrast on the transfer properties of cat retinal ganglion cells. *J. Physiol.* 285, 275–298.
- Shapley, R.M., and Victor, J.D. (1979). Nonlinear spatial summation and the contrast gain control of cat retinal ganglion cells. *J. Physiol.* 290, 141–161.
- Shapley, R.M., and Victor, J.D. (1981). How the contrast gain control modifies the frequency responses of cat retinal ganglion cells. *J. Physiol.* 318, 161–179.
- Skavenski, A.A., Hansen, R.M., Steinman, R.M., and Winterson, B.J. (1979). Quality of retinal image stabilization during small natural and artificial body rotations in man. *Vision Res.* 19, 675–683.
- Smirnakis, S.M., Berry, M.J., Warland, D.K., Bialek, W., and Meister, M. (1997). Adaptation of retinal processing to image contrast and spatial scale. *Nature* 386, 69–73.
- Srinivasan, M.V., Laughlin, S.B., and Dubs, A. (1982). Predictive coding: a fresh view of inhibition in the retina. *Proc. R. Soc. Lond. B Biol. Sci.* 216, 427–459.
- Stell, W.K., Detwiler, P.B., Wagner, H.G., and Wolbarsht, M.L. (1975). Giant retinal ganglion cells in dogfish (*Mustelus*): electrophysiology of single ON-center units. In *Vision in Fishes*, M.-A. Ali, ed. (New York: Plenum Press), pp. 99–112.
- Stone, J. (1983). *Parallel Processing in the Visual System* (New York: Plenum Press).
- Strong, S.P., Koberle, R., de Ruyter van Steveninck, R.R., and Bialek, W. (1998). Entropy and information in neural spike trains. *Phys. Rev. Lett.* 80, 197–200.
- Troy, J.B., and Enroth-Cugell, C. (1993). X and Y ganglion cells inform the cat's brain about contrast in the retinal image. *Exp. Brain Res.* 93, 383–390.
- Troy, J.B., and Lee, B.B. (1994). Steady discharges of macaque retinal ganglion cells. *Vis. Neurosci.* 11, 111–118.
- Troy, J.B., and Robson, J.G. (1992). Steady discharges of X and Y retinal ganglion cells of cat under photopic illuminance. *Vis. Neurosci.* 9, 535–553.
- van der Schaaf, A., and van Hateren, J.H. (1996). Modelling the power spectra of natural images - statistics and information. *Vision Res.* 36, 2759–2770.
- van Dijk, B.W., and Ringo, J.L. (1987). The variation in the light responses of carp retinal ganglion cells is independent of response amplitude. *Vision Res.* 27, 493–497.
- van Hateren, J.H. (1992). A theory of maximizing sensory information. *Biol. Cybern.* 68, 23–29.
- Victor, J.D. (1987). The dynamics of the cat retinal X cell centre. *J. Physiol.* 386, 219–246.
- Victor, J.D. (1988). The dynamics of the cat retinal Y cell subunit. *J. Physiol.* 405, 289–320.

Victor, J.D., and Shapley, R.M. (1979). The nonlinear pathway of Y ganglion cells in the cat retina. *J. Gen. Physiol.* *74*, 671–689.

Wandell, B.A. (1995). *Foundations of Vision* (Sunderland, MA: Sinauer).

Warland, D.K., Reinagel, P., and Meister, M. (1997). Decoding visual information from a population of retinal ganglion cells. *J. Neurophysiol.* *78*, 2336–2350.

Wässle, H., and Boycott, B.B. (1991). Functional architecture of the mammalian retina. *Physiol. Rev.* *71*, 447–480.

Winfield, D.A., and Powell, T.P. (1983). Laminar cell counts and geniculate-cortical boutons in area 17 of cat and monkey. *Brain Res.* *277*, 223–229.

Yau, K.W. (1994). Phototransduction mechanism in retinal rods and cones. *Invest. Ophthalmol. Vis. Sci.* *35*, 9–32.

Yeh, T., Lee, B.B., and Kremers, J. (1996). The time course of adaptation in macaque retinal ganglion cells. *Vision Res.* *36*, 913–931.