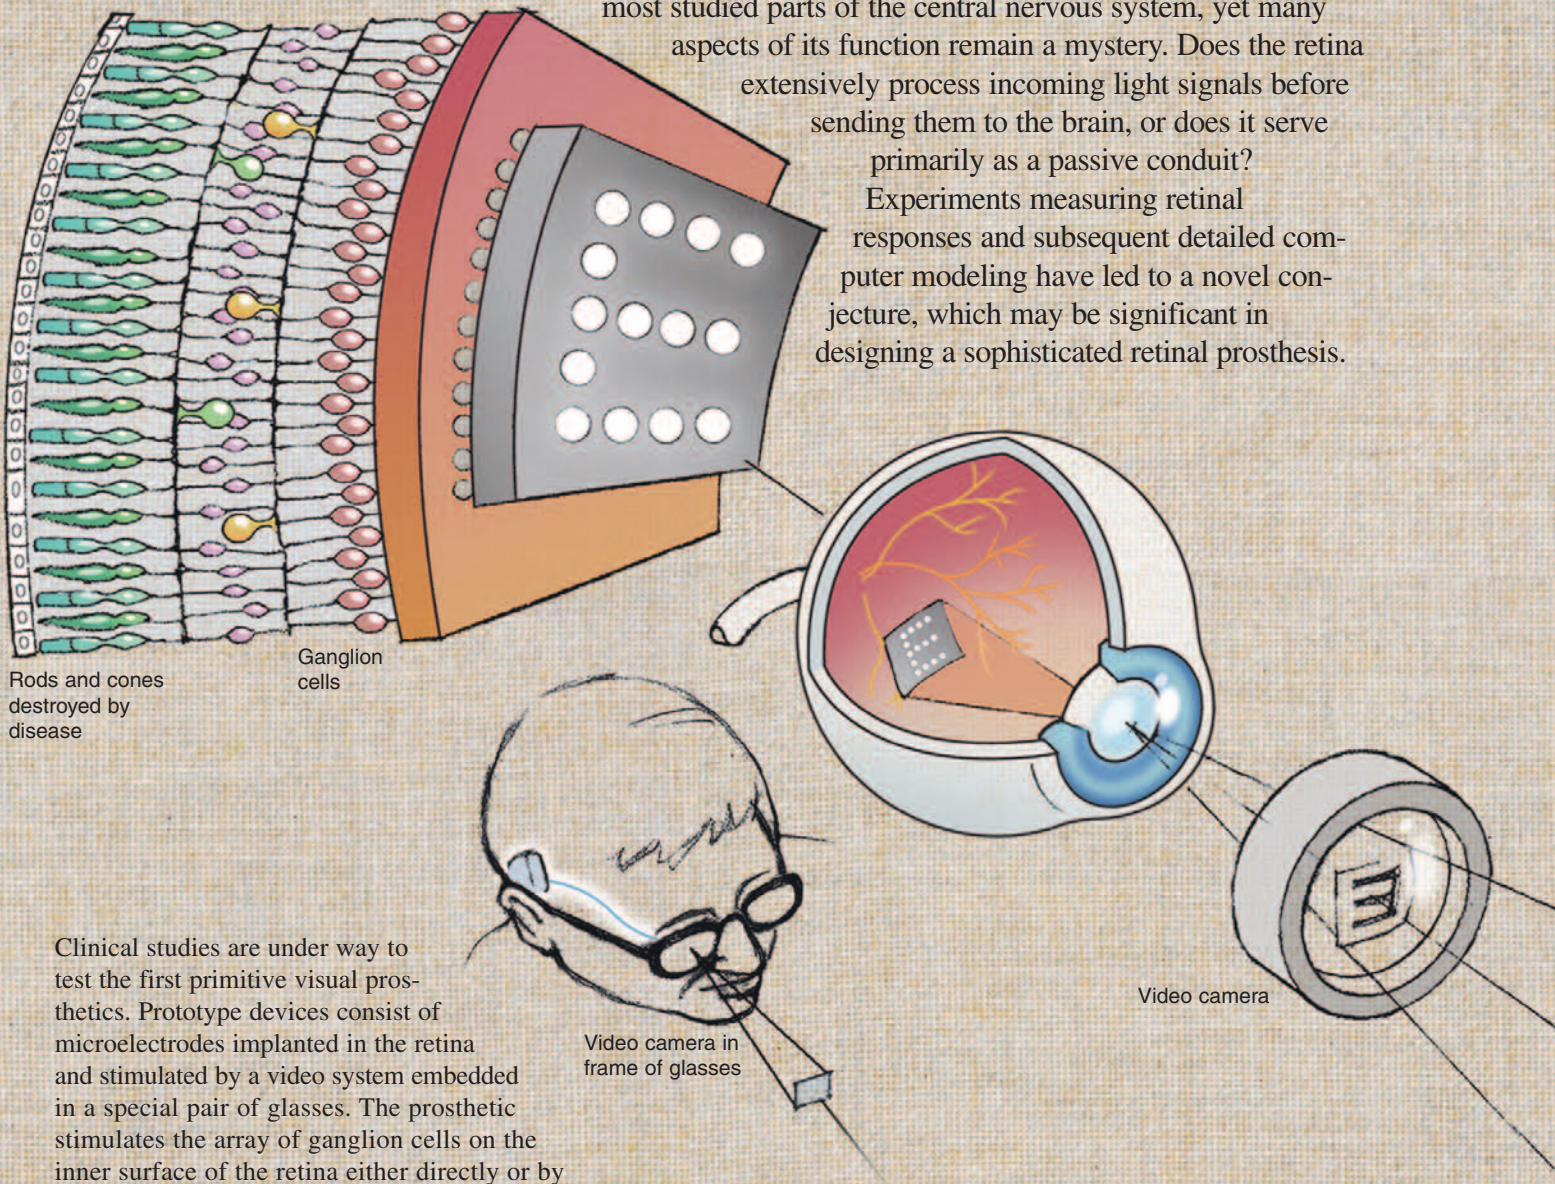


Models of the Retina with Application to the Design of a Visual Prosthesis

Garrett T. Kenyon, John George, Bryan Travis, and Krastan Blagoev

The retina, the neuronal layer at the back of the eyeball, is one of the most studied parts of the central nervous system, yet many aspects of its function remain a mystery. Does the retina extensively process incoming light signals before sending them to the brain, or does it serve primarily as a passive conduit? Experiments measuring retinal responses and subsequent detailed computer modeling have led to a novel conjecture, which may be significant in designing a sophisticated retinal prosthesis.



Rods and cones destroyed by disease

Ganglion cells

Clinical studies are under way to test the first primitive visual prosthetics. Prototype devices consist of microelectrodes implanted in the retina and stimulated by a video system embedded in a special pair of glasses. The prosthetic stimulates the array of ganglion cells on the inner surface of the retina either directly or by activating their synaptic inputs, thereby causing them to fire action potentials that propagate along the optic nerve to processing centers in the brain. Creating firing patterns that match those produced in the healthy retina by natural visual stimuli is the foremost challenge confronting the development of a retinal prosthesis.

Video camera in frame of glasses

Video camera

Some forms of adult-onset blindness are characterized by a massive loss of photoreceptors but a relative sparing of fibers in the optic nerve. Recent clinical studies suggest that patients suffering from such visual impairments could benefit from a prosthetic device capable of stimulating the remaining retinal neurons and thereby mimicking the function of the missing rods and cones. A retinal prosthesis is illustrated conceptually on the opening page of this article. The light transduction role of the damaged or missing photoreceptors is performed by a video camera attached to a pair of specially configured eyeglasses worn by the patient. The video image is processed and then transmitted, through a cable or some form of wireless telemetry, to a multielectrode array attached to the inside surface of the retina. Stimulating the multielectrode array in an approximately one-to-one spatial correspondence with the video image will hopefully produce patterns of neural activity in the optic nerve similar to those produced by an undamaged retina during normal vision. Preliminary studies, in which a crude prototype of the above design was used, yielded encouraging results (Humayun et al. 2003). Other strategies for a retinal prosthesis—particularly the insertion of a standalone photodiode array between the retina and the back of the eye—are being investigated as well. For a recent

review, see Margalit et al. (2002).


A consortium of DOE laboratories is working on a number of difficult technical problems that must be solved before a functional retinal prosthesis becomes widely available. Here, we describe computer modeling studies that have two goals regarding the optimal design of a retinal prosthesis: (1) to discover how visual information is processed and encoded by retinal circuitry, discussed in the main article, and (2) to improve our understanding of how retinal components, at the level of individual cells and across interconnected circuits, are activated by specific spatiotemporal patterns of electrical stimulation, discussed in the box, “Modeling Stimulation by a Retinal Prosthesis” on page 122. Understanding how the retina encodes visual information and how surviving elements in the diseased retina react to external stimulation is critical to achieving maximal therapeutic benefit from a prosthetic device.

Attempts to develop computer models of the retina benefit greatly from a large existing knowledge base. The anatomy and physiology of the retina have been extensively studied, especially in comparison with many other parts of the central nervous system. Moreover, the inputs and outputs of the retina have been well characterized. Because it receives no major feedback from the brain, the retina can be studied as a standalone circuit. Thus, we use experimental data to constrain our computer simulations of the retina to an extent not possible when modeling more central brain areas. Nevertheless, our studies may provide valuable insights into the design of neural prosthetics for other, less-accessible brain regions and may suggest new image-processing strategies for computer-vision systems.

The Retina

The retina consists of several layers of cells at the back of the eye that collectively are responsible for the transduction and preprocessing of visual signals (Figure 1). In photomicrographs of retinal cross-sections, several processing layers can be distinguished. In patients with certain forms of adult-onset blindness, the outermost photoreceptor layer is nearly or completely degenerated, whereas some fraction of neurons in the inner retina, particularly the ganglion cells whose axons make up the optic nerve, are spared to some extent. Such patients are potential candidates for a retinal prosthesis. However, before a prosthetic device can be optimally used, it is vital to understand how visual information is encoded in the pattern of electrical impulses traveling down the optic nerve.

The output of the retina, and indeed of most neurons in the central nervous system, cannot be classified in conventional electrical engineering terms as either analog or digital; rather, neuronal output consists of a temporal sequence of impulses, or spikes (see Figure 2). It is therefore vital to understand how visual information is encoded in spike trains traveling down the optic nerve. In the absence of stimulation, most ganglion cells fire spikes randomly at a background rate much lower than their maximum firing frequency. The conceptual diagram in Figure 2 depicts the spike trains from two clusters of neighboring ganglion cells. Clusters are indicated rather than single ganglion cells both to justify the high signal-to-noise ratio illustrated in the figure and to take into account the spatial convergence of optic-nerve fibers onto target neurons deep within the brain. When a cluster is stimulated by a small spot roughly equal in size to the excitatory portion of the receptive field (the local region of the visual






Figure 1. The Retina

(a) Located at the back of the eyeball, the retina consists of many types of neurons arranged in a layered structure. The light-sensitive cells (photoreceptors—rods and cones) are in the outermost layer, farthest from the incoming light. In front of the photoreceptors are neurons that perform specialized processing. At the innermost layer are ganglion cells whose axons make up the optic nerve. (b) A photomicrograph of a cross section of the retina reveals those distinct processing layers.

(Courtesy of Webvision <http://webvision.med.utah.edu/> at the Moran Eye Center.)

space to which the indicated group of cells best respond), the firing rate, that is, the average number of spikes per time interval, increases markedly in proportion to the contrast between the spot's intensity and the light intensity immediately surrounding the spot. Regardless of the stimulus intensity, however, the timing of the individual spikes in response to a small spot remains more-or-less randomly distributed such that the spikes elicited by two small spots will typically be entirely uncorrelated. The observation that the mean firing rate is proportional to the stimulus intensity while the spikes themselves are distributed randomly in time is the basis for the rate-code hypothesis, which posits that information is transmitted only by the mean number of spikes per time interval irrespective of their precise timing.

There is evidence, however, that the rate-code hypothesis is incomplete. As the size of the two spots increases, the total number of spikes per time interval is reduced somewhat

by lateral inhibition, but the most striking effect is the appearance of an oscillation in the firing rate, causing spikes to occur in relatively narrow clusters, or bursts. The phase of the underlying oscillation drifts randomly over time, so that the bursts evoked by separate spots will rapidly become uncorrelated even if both sets of neurons are modulated at a similar frequency. Remarkably, when the two groups of neurons are stimulated by a single large spot, the groups' underlying oscillations become strongly phase-locked, suggesting that the relative timing of spikes in the optic nerve can convey information about the spatial connections of features in the visual field. Such information is not conveyed by the local firing rate. To better illustrate the above encoding principles, it is useful to examine real physiological data. Figure 3 shows an intracellular recording from a retinal ganglion cell in response to a sinusoidally varying light intensity. In Figure 3(a), the intensity of the light is shown as a function of time.

Conceptually, the recorded trace in Figure 3(b) can be divided into two parts: a subthreshold membrane potential, exhibiting an approximately sinusoidal modulation, and action potentials, or spikes, which are the large impulses riding on top of the subthreshold membrane potential. This potential is not available to the brain because it represents the analog component of the response that is most directly proportional to the incident light intensity. Only the spikes riding on top of this potential are transmitted through the optic nerve to relay nuclei within the brain. Because each spike is, to a first approximation, identical to every other spike, information can be conveyed only by the temporal sequence of the spikes. To reveal the information embedded in experimentally recorded spike trains, neuroscientists typically average over many stimulus trials. The response histogram, obtained by combining spike trains from many stimulus trials, shows that the average firing rate of the recorded ganglion cell is roughly

proportional to the applied light intensity, except for an approximately 90° phase advance, which reflects the fact that the cell is sensitive to the rate of change of the light intensity, and a negative cutoff due to the fact that the firing rate cannot drop below zero—see Figure 3(c). According to the rate-code hypothesis, the multitrial rate histogram fully characterizes the information conveyed by neural spike trains. (Of course, the brain cannot perform a multitrial average, but it is assumed that the brain can extract similar information in real time by combining low-pass filtering and information from many cells.)

About 10 years ago, Wolf Singer's laboratory in Germany reported that retinal neurons also use the relative timing of spikes to encode information about visual stimuli that is not conveyed by their local firing rates (Neuenschwander and Singer 1996, Neuenschwander et al. 1999). Unlike the previous example, which involved an intracellular recording from a single cell, the data from Singer's laboratory consists of spike trains from two retinal neurons recorded simultaneously in response to a sustained light stimulus, either two separate short bars or a single long bar (Figure 4). In both cases, the edges of the bar stimuli were positioned over the receptive centers of the recorded cells so that, locally, the stimulation was approximately the same regardless of whether the stimulus consisted of one or two bars. Another difference from the previous example is that the spike trains in the Singer experiment were recorded with extracellular electrodes, which do not permit examination of the corresponding membrane potentials. However, spike trains recorded immediately outside the cell can be analyzed for temporal structure both individually and by pairs.

Correlation functions constructed from the individual spike trains revealed that the outputs of both reti-

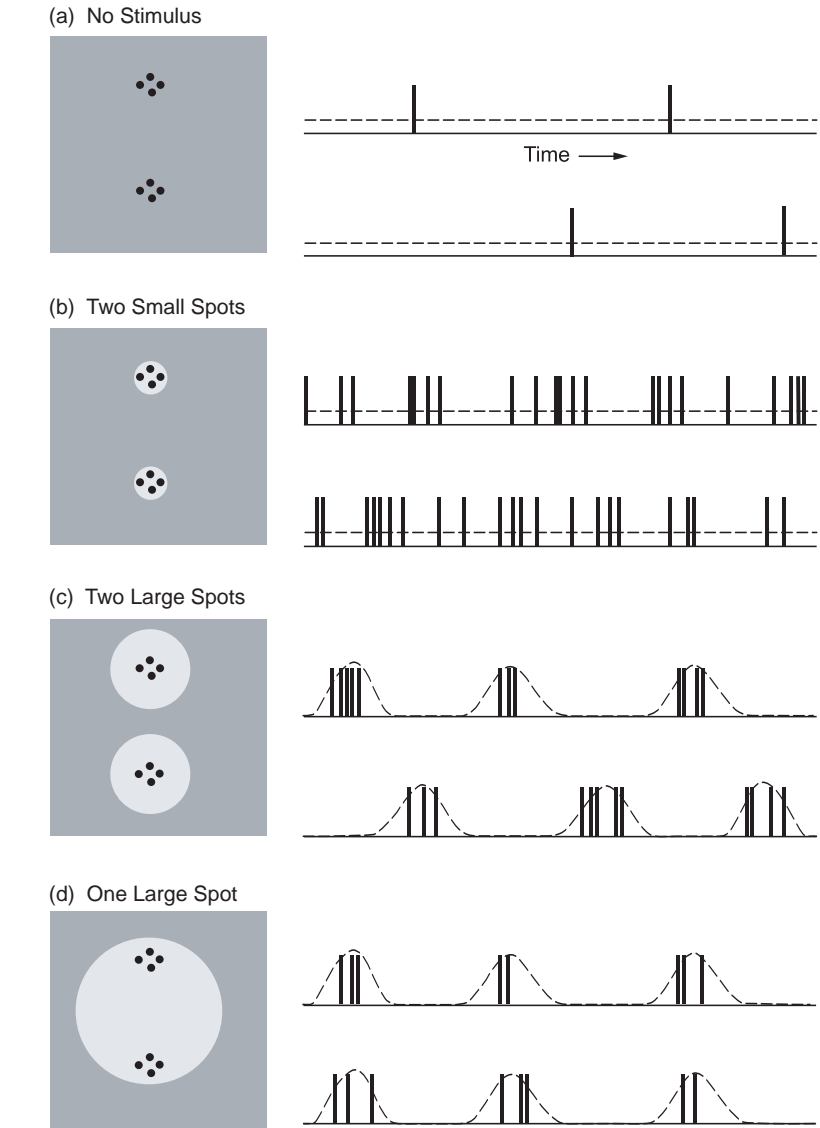


Figure 2. Retinal Responses to Light Stimuli

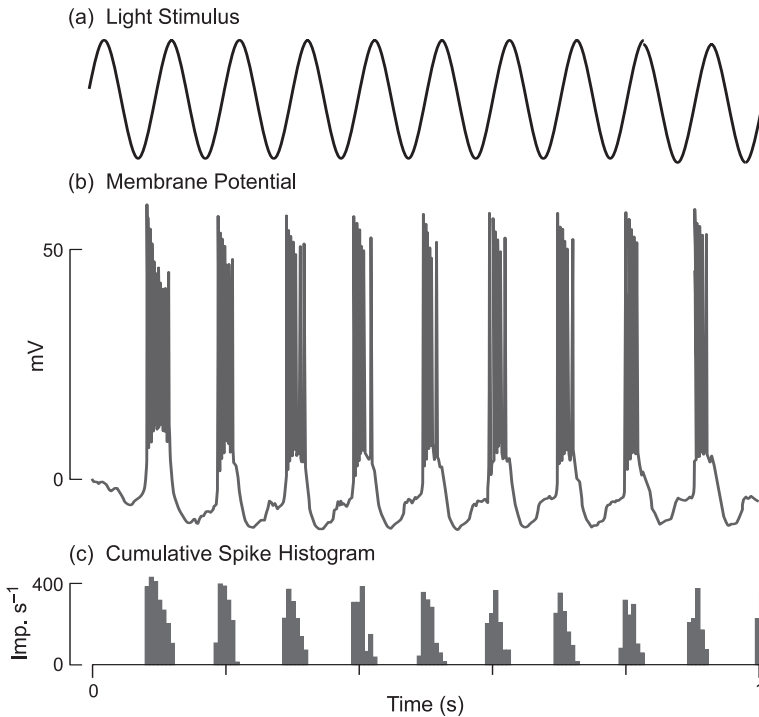
The gray boxes show two groups of retinal photoreceptors (black dots) exposed to spots of light during four experiments. At right are the outputs of the two groups of ganglion cells activated by the two groups of photoreceptors. (a) With no light stimulus, the outputs are random spikes. (b) When each group of photoreceptors is exposed to a small spot of light, the average firing rate of the overlying ganglion cells increases in proportion to the ratio of the intensity of the spot to the intensity of the region immediately surrounding the spot, but the spikes still occur more or less randomly. (c) As the spot size increases, the firing rate decreases somewhat, and the spikes bunch up. However, the bunches are not synchronized. (d) When the spot size is large enough to cover both groups of photoreceptors, the bunches become synchronized.

nal neurons were modulated by periodic oscillations—see the small boxes in Figures 4(a) and 4(b). The oscillations elicited by separate light


bars were not phase-locked, as indicated by the relative absence of correlations between the two spike trains when stimulated by separate bars—

Figure 3. Responses of a Single Retinal Ganglion Cell

A motion-sensitive ganglion cell from a mammalian retina was stimulated by a spot of light whose intensity varied sinusoidally in time. (a) The intensity of the light spot is shown as a function of time. (b) The subthreshold membrane potential recorded at the soma (cell body) consists of action potentials, or spikes, riding on top of an approximately sinusoidal modulation. Only the spikes are transmitted to relay nuclei in the brain. (c) A response histogram is constructed from spikes accumulated over many identical stimulus trials, showing that the average firing rate is approximately proportional to the applied stimulus intensity except for a phase shift and a lower cutoff at 0 Hz. (Dacey and Lee 1994. Reprinted with permission from *Nature*.)



(a) Two Separated Light Stimuli



(b) A Single Light Stimulus

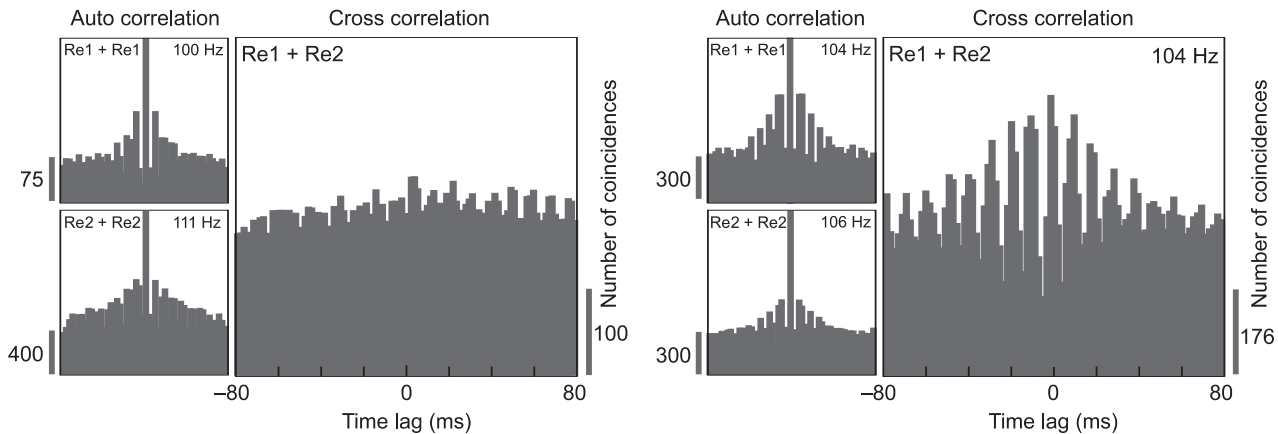



Figure 4. Responses of Two Retinal Ganglion Cells

Two cat-retina ganglion cells separated by 6 degrees are monitored by electrodes Re1 and Re2, respectively. Spike trains generated in response to rectangular light stimuli were recorded simultaneously from each electrode, and autocorrelation and cross-correlation histograms were computed. (a) When two distinct rectangular light stimuli are used, strong oscillations are present in the autocorrelation histogram of each ganglion cell, but the cross-correlation histogram between the two ganglion cells is essentially flat. The sharp peaks in the autocorrelation histograms correspond to the spike bunching in Figures 2(c) and 2(d). (b) When a single rectangular stimulus that connects the region between the cells is used, the oscillations in firing rate of the two ganglion cells become strongly phase locked, generating a strong oscillation in the cross-correlation histogram. The cross-correlation oscillation corresponds to the synchronization of the spike bunches in Figure 2(d). (Neuenschwander 1996. This figure was redrawn courtesy of *Nature*.)

see the large box in Figure 4(a). On the other hand, the evoked oscillations were tightly phase-locked when activated by a single bar, as indicated by the strong periodic modulations in the corresponding correlation function illustrated in the large box in Figure 4(b). Thus, the timing of retinal ganglion-cell spikes, especially with respect to the phase differences between the oscillatory responses of separate groups of ganglion cells, can convey information relevant to the spatial separation or spatial binding of visual features.

A Computer Model

We constructed a computer model of the retina (Figure 5) to study how temporal codes in the retina might be generated and what types of visual information such codes might convey. Very few circuits in the central nervous system have been completely characterized, and the circuits of the retina are no exception. However, enough is known about retinal anatomy and physiology to allow the construction of a model that accounts for many aspects of experimentally recorded light responses in a manner consistent with general patterns of neuronal connectivity found in most vertebrate species. Moreover, by requiring the model to account for a wide range of experimental data, we were able to infer some aspects of the unknown anatomy and physiology. In particular, by increasing the range of observed phenomena explained by the model, we were able to eliminate alternative implementations that were inconsistent with published recordings.

The model consisted of five distinct cell types illustrated in Figure 5(a): bipolar cells, three classes of amacrine cells, corresponding to the small, large, and polyaxonal subtypes, and ganglion cells. Bipolar cells are




Figure 5. Computer Model of the Retina

(a) Our computer model consisted of five cell types: bipolar (BP) cells, small (SA), large (LA), and polyaxonal (PA) amacrine cells, and alpha ganglion (GC) cells, arranged in a 32×32 square mosaic with wrap-around boundary conditions.

Although this side view of one of the mosaic's units shows only two BPs, there were actually four BPs. Light stimuli were simulated by injecting currents directly into the BPs. (Photoreceptors were not included in the model.) The inhibitory connections can be organized into three categories: *Feedforward and feedback inhibition*. Excitatory synapses from BPs were balanced by a combination of reciprocal synapses and direct inhibition of the GCs, mediated by the nonspiking amacrine-cell types. *Serial inhibition*. The three amacrine-cell types regulated each other through negative feedback loops. *Resonance circuit*. The PAs were excited locally through electrical synapses with GCs, and their axons gave rise to widely distributed inhibition that contacted all cell types, but most strongly the GCs and other PAs. Not all connections present in the model are shown. (b) A simplified schematic diagram of the computer model shows how a combination of local excitation (triangles) carried by gap junctions (resistors) and long-range inhibition (empty circles) carried through axon-bearing amacrine cells (orange dotted lines and filled black circles) produced physiologically realistic oscillations dependent on stimulus size.

(Reprinted with the permission of Cambridge University Press.)




Figure 6. Light-Evoked Oscillations in the Retinal Model
(a) The average correlation function of the retinal model's output for pairs of ganglion cells exhibits an oscillation whose amplitude, frequency, and duration are similar to those of **(b)**, the correlation function for experimentally recorded spike trains from cat retina cells in response to an analog stimulus.

relay neurons that receive synaptic input from photoreceptors (not shown) and provide the principal excitatory drive to the other neuron types. The bipolar cells are therefore critical elements in the “vertical” pathway representing the main direction of information flow from the photoreceptors to the optic nerve. On the other hand, amacrine cells mediate lateral interactions that are essential for processing and encoding visual signals. Indeed, without amacrine-cell interactions, corresponding to the “horizontal” pathway, the output of the retina would simply replicate the activity across the photoreceptor array. Nearly 30 different kinds of

amacrine cells have been described, and it would not have been possible to include them all in the model. Instead, we included only the minimum number of amacrine-cell types necessary to account for the basic features of retinal light responses, as well as for the synchronous oscillations evoked by large stimuli.

While the circuitry incorporated into the model is somewhat complicated, the connections can be grouped into three major categories:

Excitation. The bipolar cells, which relayed visual signals from the photoreceptors to all the other cell types, were the only source of excitation in the model.

Local inhibition. The model amacrine cells made feedforward inhibitory synapses onto the ganglion cells and feedback inhibitory synapses onto the bipolar cells and serially inhibited each other. These local inhibitory synapses acted to increase the dynamical range of the model retina, by negative feedback, and further contributed to shaping light-evoked activity so as to amplify the responses to both spatial and temporal contrast.

Long-range axon-mediated feedback. The polyaxonal amacrine cells received local excitation from ganglion cells by electrical synapses or gap junctions and, in turn, made long-range inhibitory connections onto all cell types. This delayed negative feedback circuit accounted for the generation of oscillatory responses in the model retina and has been redrawn in Figure 5(b) to emphasize the circuit's essential components and their interconnections. Further details of the model, particularly its ability to account for experimental data as well as its numerical stability and robustness to parameter variation, are available elsewhere (Kenyon et al. 2004a, Kenyon et al. 2004b, Kenyon et al. 2003a, Kenyon et al. 2003b).

Oscillations

To assess the stimulus-evoked oscillations in the retinal model, correlations were computed between the spike trains arising from all pairs of ganglion cells activated by a large spot, and the results were combined into an averaged correlation measure—refer to Figure 6(a). The amplitude, frequency, and persistence of the periodic modulations in the averaged correlation function obtained from the retinal model were qualitatively similar to those observed experimentally in recordings of cat ganglion cells responding to large, high-contrast spots, as illustrated in Figure 6(b). For both the cat retina and the retinal model, the correlation amplitude falls off with increasing delay, eventually returning to approximately baseline levels after several cycles of the underlying oscillation. In both sets of data, the phases of the underlying oscillations drift randomly over time, so that firing activity becomes uncorrelated over sufficiently long delays. This time drift is a fundamentally nonlinear phenomenon arising from the threshold nature of spike generation. In contrast, the phase of a linear harmonic oscillator is always fixed relative to the stimulus onset. The retinal model thus captures an essential nonlinear property of the biological circuitry. Moreover, the good qualitative agreement between theory and experiment implies that the free parameters in the model, particularly those involving the axon-mediated feedback circuit, were likely to be reasonably close to their true physiological values. Other comparisons with physiological data were used to verify additional aspects of the model.

The retinal model was also able to account for the experimentally observed size dependence of retinal oscillations (Figure 7). Specifically, the oscillations evoked by stimuli of various sizes in our retinal model

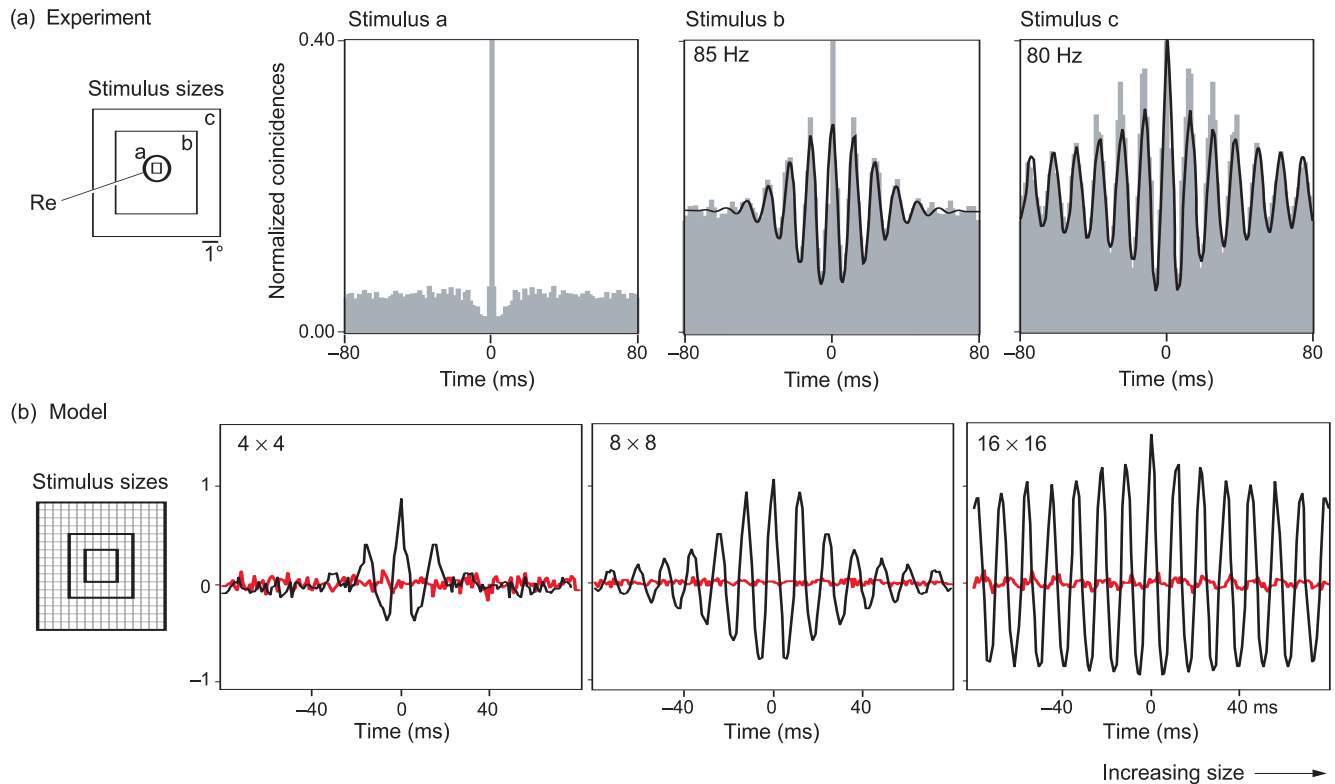


Figure 7. Stimulus-Size Dependence of Retinal Oscillations

(a) The correlation function computed between experimentally recorded spike trains from cat retina cells exhibits a strong increase in oscillatory activity with increasing stimulus size. (b) The correlation function computed for the oscillations produced by our retinal model exhibits a similar size dependence. (Experimental data from Neuenschwander (1996). Redrawn courtesy of *Nature*.)

were similar to those measured from the cat retina. In both sets of data, small stimuli evoked little or no oscillatory response, whereas large stimuli evoked oscillations with very large amplitudes. Because the axon-mediated feedback was spread out over a wide retinal area, only large stimuli could evoke strong oscillatory responses. The notion that oscillations are associated with large stimuli led us to put forward a novel hypothesis about the types of visual signals encoded in the periodic temporal structure of retinal spike trains. We discuss this hypothesis in more detail below.

The model also accounted for a high-frequency resonance observed in the responses of certain retinal neurons to temporally modulated stimuli (Figure 8). The temporal modulation transfer function (tMTF) measures

how strongly the output of a system is modulated as a function of the frequency of a sinusoidal input.

Harmonic or oscillatory systems typically exhibit a resonance frequency, at which the output of the system can be driven to relatively large amplitudes. The presence of high-frequency oscillations in retinal light responses suggests that there will be a corresponding resonance in the tMTF, given by the amplitude of the sinusoidal modulation in the ganglion-cell firing rate when plotted as a function of the frequency of the applied stimulus. As expected, both the cat and model retinas exhibit sharp resonance peaks in their tMTFs at frequencies above 60 hertz, at which frequency value oscillatory responses are also observed. Moreover, the model accounted not only for the resonance itself but also

for the associated kink in the phase-response curve, which plots how much the phase of the output modulation is retarded or advanced relative to the sinusoidal input. Such kinks are not present in the phase-response curves of simple harmonic oscillators. Using our computer model, we were able to show that the kink in the phase-response curve obtained from retinal ganglion cells was due to entrainment. When driven at relatively low modulation frequencies, the oscillations produced by retinal circuitry, whose phase drifts randomly over time, quickly become independent of the phase of the driving stimulus. As the frequency of the driving stimulus approaches the resonance frequency, however, the two oscillations become entrained, causing an abrupt advance in the phase-response curve. Such resonances may

Figure 8. Temporal Modulation Transfer Functions (tMTFs)

(a) The tMTF recorded from cat-retina ganglion cells is obtained by plotting the magnitude of the fundamental Fourier component in the response histogram as a function of the temporal modulation frequency of the applied stimulus. The maximum response occurs at a broad low-frequency resonance, between 10 Hz and 20 Hz, but there is also a prominent high-frequency resonance at around 70 Hz. (b) The phase-response curve corresponding to the data shown in (a) is plotted as a function of temporal modulation frequency. The phase-response curve exhibits a prominent kink at frequencies near the rising phase of the resonance peak. (c) A similar resonance is present in the tMTF recorded from ganglion cells in the retinal model in response to a temporally modulated spot. (d) Likewise, the phase-response curve of the model ganglion cells also exhibits a kink near the onset of the resonance peak. The kink is caused by entrainment of the retinal oscillations by the applied stimulus.

(Reproduced from the *Journal of General Physiology*, 1987, Vol. 89, pp. 599–628, by copyright permission of The Rockefeller University Press.)





Figure 9. Stimulus-Selective Oscillations

(a) Two bar-shaped light stimuli are shown in relation to the receptive field centers of four simultaneously recorded ganglion cells. Cross-correlation histograms were computed during the plateau portion of the responses for pairs of ganglion cells at opposite ends of the same bar or at opposing tips of separate bars. All ganglion cell pairs were separated by 7 diameters. The cross-correlation histograms were computed for pair 1,2 from the upper bar (b), pair 2,3 from the two separate bars (c), and pair 3,4 from the lower bar (d). The histograms exhibit significant oscillations only for pairs stimulated by the same bar.

(Courtesy of Wolf Singer and colleagues.)



also be relevant to the effective operation of a retinal prosthesis by enabling strategies for selectively activating certain types of retinal neurons at the characteristic frequencies of stimulation.

Finally, the retinal model was able to reproduce the stimulus selectivity of retinal oscillations first reported by Wolf Singer's laboratory, as outlined above. By examining the relative timing of spikes produced by retinal ganglion cells responding to either the same or to different objects, we were able to show that model elements activated by the same large object were strongly correlated, or phase locked, by a common underlying oscillation at a frequency of approximately 100 hertz (see Figure 9). Pairs of model retinal neurons activated by different objects, however, were not correlated; that is, the phases of their underlying oscillations varied randomly with respect to each other. Thus, our retinal model captures the interesting property of biological neurons that their evoked oscillations in responses to large visual features are stimulus specific and are only phase locked for cells responding to the same contiguous object. The above results illustrate how the retinal model was able to account for many of the main experimentally observed aspects of oscillatory phenomena.

What Do Oscillations Encode?

Having established the biological plausibility of the retinal model, we then used computer simulations to explore what information stimulus-specific oscillations might convey to the brain. Because our model is consistent with the known anatomy and physiology of the cat retina, the model can provide a useful tool for investigating how information can be encoded in the temporal structure of spike

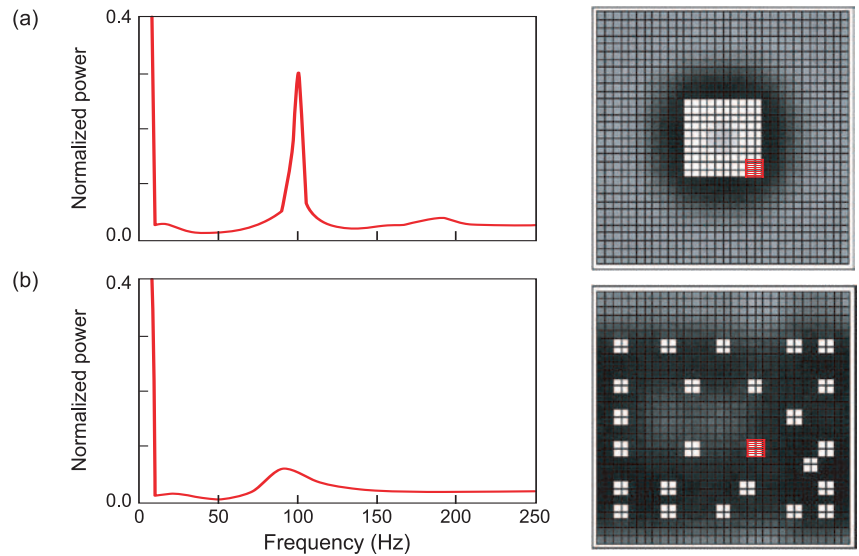


Figure 10. See Globally, Spike Locally

We exposed 25 clusters of ganglion cells in our computer-model array of 32×32 cells to two different high-contrast light stimuli and computed the power spectrum from the output of a single cluster (shown in red) for each exposure condition. Each cluster consisted of 2×2 neighboring ganglion cells. The power spectra (left) were normalized by the total average firing rate for each exposure condition (right). (a) For a cluster that was part of a large illuminated area, the power spectrum peaked sharply between 60 Hz and 120 Hz. (b) For a cluster illuminated in isolation, the power spectrum exhibited only a small hump.

trains propagating down the optic nerve. Based on the stimulus selectivity of retinal oscillations, as well as their observed size dependence, a collection of disconnected spots will elicit only weak periodic modulations in optic-nerve fibers, whereas a single large stimulus will elicit strong period modulations. To test this idea, we exposed our computer-model array of 32×32 ganglion cells to two different light stimuli. The first stimulus was a large, square spot covering 25 contiguous clusters of cells—refer to Figure 10(a). Each cluster consisted of 2×2 cells. For this stimulus, the power spectrum of the spike trains from a single cluster exhibited a large, sharp peak at around 100 hertz. However, when the array was exposed to 25 small, isolated spots, each of which covered exactly one cluster but otherwise elicited approximately the

same total number of spikes per time interval, there was only a small hump in the power spectrum, as shown in Figure 10(b). These results suggest that the periodic temporal structure in the spike trains obtained from small clusters of neighboring neurons encodes the overall size of the object to which the clusters respond.

To investigate the above hypothesis, we were guided by two principles: (1) Because it takes us only a fraction of a second to form a visual impression, the information conveyed by stimulus-specific oscillations must be available on short, physiologically meaningful time scales—roughly a few hundred milliseconds. (2) Because the spatial convergence of retinal neurons onto target cells in the brain is rather low, with each target cell receiving input from only a few retinal neurons, the




Figure 11. Size Discrimination

A Bayes discriminator was used to classify light spots as either “smaller” or “larger” from the single-trial oscillatory activity of (a) cat retina ganglia or (b) the ganglia in our retinal computer model. For the multiunit spike trains recorded from cat retina ganglia, the percentage of correctly classified trials ranged from ~73% to ~87% as the length of the multiunit spike train segment increased from 50 to 300 ms. The percentage of correctly classified trials using multiunit spike trains from the retinal model improved with longer analysis windows and as more ganglion cells were included in the spike record.

information conveyed by stimulus-specific oscillations must be available locally in the firing activity of a similar number of neighboring cells. We therefore used the retinal model to quantify the information conveyed about the global properties of a stimulus, in this case the total size of the object, by a 2×2 neighborhood of retinal output neurons in a few hundred milliseconds. At the same time, having received data from Wolf Singer’s laboratory recorded from output neurons in the cat retina under analogous experimental conditions, we were able to test directly the predictions of our retinal model.

One of our studies tested our ability to determine if a group of neighboring cells was responding to a small or a large object from the group’s local firing activity alone. In Figure 11, we plot the results of this study in terms of the “accuracy of size discrimination,” which was equal to the fraction of trials in which the total size of the

stimulus could be correctly inferred from the local firing activity. Random events were added to the model spike trains to ensure that the firing rate did not change as a function of stimulus size. The only cue available from the local firing activity regarding the total size of the stimulus was therefore the amplitude of the synchronous oscillations. Our results showed that, in model spike trains 300 milliseconds long, using as few as four spike trains from a small neighborhood (2×2 cells), nearly perfect accuracy can be achieved. The accuracy for the experimentally recorded spike trains was slightly lower, possibly reflecting to some extent the suboptimal recording conditions in which several different ganglion cell types contributed to the multiunit response. Overall, our modeling results imply that there is a tradeoff between the number of cells included in the analysis and the total time allowed for accomplishing the size-discrimination task. Specifically,

as more cells were included in the analysis, shorter periods were required to achieve the same accuracy.

Why would it be important for retinal neurons to convey information about stimulus size in their local firing activity? Studies of a frog’s retina may provide the answer. Tachibana’s laboratory in Japan has shown that the frog retina has specialized neurons, called dimming detectors, that exhibit strong synchronous oscillations when activated by a large dimming object but do not exhibit such oscillations when activated by a small dimming object (Ishikane et al. 1999). To a frog, a small dimming spot could be a fly or other food source, but a large dimming spot is more likely to be a bird or other dangerous predator. In this situation, one can easily appreciate why size matters. ■

Acknowledgments

We gratefully acknowledge the following collaborators in undertaking the described studies: Greg Stephens, Mark Flynn, and Benjamin Barrowe.

Further Reading

- Frishman, L. J., A. W. Freeman, J. B. Troy, D. E. Schweitzer-Tong, and C. Enroth-Cugell. 1987. Spatiotemporal Frequency Responses of Cat Retinal Ganglion Cells. *J. Gen. Physiol.* **89**: 599.
- Geddes, L. A., and L. E. Baker. 1967. The Specific Resistance of Biological Material—A Compendium of Data for the Biomedical Engineer and Physiologist. *Med. Biol. Eng.* **5** (3): 271.
- Humayun, M. S., J. D. Weiland, G. Y. Fujii, R. Greenberg, R. Williamson, J. Little, et al. 2003. Visual Perception in a Blind Subject with a Chronic Microelectronic Retinal Prosthesis. *Vision Res.* **43**: 2573.
- Ishikane, H., A. Kawana, and M. Tachibana. 1999. Short- and Long-Range Synchronous Activities in Dimming Detectors of the Frog Retina. *Vis. Neurosci.* **16** (6): 1001.

- Jensen, R. J., J. F. Rizzo III, O. R. Ziv, A. Grumet, and J. Wyatt. 2003. Thresholds for Activation of Rabbit Retinal Ganglion Cells with an Ultrafine, Extracellular Microelectrode. *Invest. Ophthalmol. Vis. Sci.* **44** (8): 3533.
- Kenyon, G. T., J. Theiler, J. S. George, B. J. Travis, and D. W. Marshak. 2004. Correlated Firing Improves Stimulus Discrimination in a Retina Model. *Neural Comput.* **16** (11): 2261.
- Kenyon, G. T., J. Theiler, D. W. Marshak, B. Moore, J. Jeffs, and B. J. Travis. 2003. Firing Correlations Improve Detection of Moving Bars. In *Proc. Int. Joint Conf. Neural Network.* **1-4**: 1274.
- Kenyon, G. T., B. J. Travis, J. Theiler, J. S. George, G. J. Stephens, and D. W. Marshak. 2004. Stimulus-Specific Oscillations in a Retinal Model. *IEEE Trans. Neural Network.* **15** (5): 1083.
- Kenyon, G. T., B. Moore, J. Jeffs, K. S. Denning, G. J. Stephens, B. J. Travis et al. 2003. A Model of High-Frequency Oscillatory Potentials in Retinal Ganglion Cells. *Vis. Neurosci.* **20** (5): 465.
- Margalit, E., M. Maia, J. D. Weiland, R. J. Greenberg, G. Y. Fujii, G. Torres et al. 2002. Retinal Prosthesis for the Blind. *Surv. Ophthalmol.* **47**(4): 335.
- Neuenschwander, S., M. Castelo-Branco, and W. Singer. 1999. Synchronous Oscillations in the Cat Retina. *Vision Res.* **39**: 2485.
- Neuenschwander, S., and W. Singer. 1996. Long-Range Synchronization of Oscillatory Light Responses in the Cat Retina and Lateral Geniculate Nucleus. *Nature* **379**: 728.
- Travis, B. J., and A. D. Chave. 1989. A Moving Finite Element Method for Magnetotelluric Modeling. *Phys. Earth Planet. Inter.* **53**: 432.
- Unsworth, M. J., B. J. Travis, and A. D. Chave. 1993. Electromagnetic Induction by a Finite Electric Dipole Source over a 2-D Earth. *Geophys.* **58** (2): 198.

*For further information, contact
Garrett Kenyon (505) 667-1900
(gkenyon@lanl.gov).*

Modeling Stimulation by a Retinal Prosthesis

We have begun developing a three-dimensional (3-D) model of the retinal extracellular space. Existing software, originally developed for modeling the flow of ground water through porous material [(Unsworth et al. 1989), (Travis and Chave 1989)], has been adapted to calculate the potential gradients produced by an arbitrary distribution of stimulating electrodes. The flow of water in porous media (that is, sedimentary rock) and the flow of current through the extracellular space are mathematically identical problems, allowing powerful software tools developed in one context to be applied to the other. Our basic strategy is to avoid the fully interacting problem that requires solving for the intercellular and extracellular potentials simultaneously. Instead, we take advantage of the fact that the ephaptic, or incidental coupling between retinal neurons via the extracellular space is small and that any significant extracellular potential gradients will be due almost entirely to external stimulation.

Prosthetic stimulation can therefore be modeled in two distinct steps: (1) calculate the extracellular potentials produced by the applied currents and (2) compute how the resulting gradients act upon dendritic and axonal processes within the retina. In principle, the same technology could be used in reverse; the normal light-evoked responses of retinal neurons could be calculated before hand and the resulting membrane currents could be used as sources to estimate local field potentials. Such technology could eventually be useful for connecting realistic simulations of retinal circuits to clinical measures such as the electroretinogram.

As preliminary data, we have developed a simplified model of the retina and associated structures in which the various anatomical ele-

ments, consisting of the vitreous, retina and retina pigment epithelium (RPE)/choroid, as well as the multielectrode array itself, were represented as rectangular blocks (Figure A). The bulk conductivity of each element was based on published values (Geddes and Baker 1967), with the bulk conductivity of the multielectrode array set to zero. The computer model was used to calculate the extracellular potential gradients produced by a 1 microampere anodic current pulse passed through a single stimulating electrode, 10 microns in diameter, located on the vitreous surface. The return electrode was placed in the vitreous cavity 400 microns away along a line perpendicular to the retinal surface. It is well established that

cathodic current pulses are much more effective for stimulating retinal neurons, but for modeling the spatial distribution of extracellular currents, the overall sign is irrelevant.

In the absence of an insulating barrier above the retina, transverse bipolar stimulation produced a dipole field that was nearly mirror symmetric, with the slight deviations arising from the conductivity differences of the various tissue components—Figure A(1). The spatial profiles of the extracellular potentials parallel to the retinal surface were examined as a function of depth from the stimulating electrode—Figure A(2). At a depth of 50 microns, the maximum value of the extracellular potential directly underneath the electrode was just under 6 millivolts,




Figure A. The Influence of Both Anatomical and Non-Anatomical Structures of the Distribution of Extracellular Currents

(1) This contour plot is of extracellular potentials due to dipole stimulation of the retina and associated structures (see labels). (2) Profiles of extracellular potentials are shown at three different depths, 50, 100, and 150 μ . Panels (3) and (4) show the same organization as (1) and (2) with the addition of an insulating block representing the prosthetic multielectrode array itself.

and fell off laterally with a length constant on the order of 100 microns. Deeper in the retina, the extracellular potential fell off more gradually as the radial component away from the electrode became smaller in the lateral direction.

A prosthetic device would not consist of a single pair of electrodes, however, but of a multielectrode array contained in an insulating package. We therefore used the computer model to examine how a large insulator affixed to the vitreous surface would affect current flow within the retina—Figures A(3) and A(4). Our results show that by forcing more of the current into the retina, a large non-conducting barrier can substantially enhance the effects of prosthetic stimulation. Inserting a representation of the prosthetic device into the 3-D model approximately doubled the extracellular potential gradients produced by the same 1 microampere current pulse applied previously. These results illustrate the general principle of how the 3-D geometry of the retina and associated structures, as well design of the prosthetic device itself, can have a large impact on the spatial distribution of externally applied currents.

Computer models can also be used to investigate how cellular properties, such as dendritic morphology and orientation, influence responses to prosthetic stimulation. As a preliminary step, we examined the changes in membrane potential produced by a 1-microampere cathodic current as a function of orientation, either vertical or horizontal—Figure B. A 50-micron passive segment, representing a bipolar cell axon or ganglion cell dendrite, was centered 75 microns from the vitreous surface directly underneath the stimulating electrode. When the passive cable segment was oriented vertically, as would likely be the case for bipolar cell axons, there was approximately a 1 millivolt depo-




Figure B. Effects of Orientation
Shown here is the change in membrane potential along a 50- μ passive cable in response to a 1- μ A current. A vertically oriented segment experiences a maximum depolarization of nearly 1 mV at its proximal tip, while a horizontally oriented segment is virtually unaffected by the stimulus. Cable centers were 75 μ from the electrode.

larization at the proximal tip closest to the electrode. On the other hand, when the same passive segment was oriented horizontally, as would be predominantly the case for ON ganglion cell dendrites, there was virtually no change in the membrane potential.

The strong influence of orientation is a direct consequence of the fact that neural processes are activated by gradients in the extracellular potential and are insensitive to the average magnitude, or constant offset. For transverse stimulation in which the dipole axis is perpendicular to the retinal surface, vertically oriented processes lay across

equipotential contour lines and thus along the maximum gradient. Horizontally oriented processes, on the other hand, lie mostly parallel to equipotential contour lines and thus experience minimal gradients along their length. The dendrites of ON ganglion cells tend to be oriented laterally and are therefore likely to have higher activation thresholds than vertically oriented processes. Ganglion cell axons are also oriented in a predominantly lateral direction and thus are less strongly activated by transverse currents, but this effect is potentially countered by their closer proximity to the electrode array.

Finally, if an anodic current pulse had been applied instead, the proximal tip of the vertically oriented segment would have been hyperpolarized rather than depolarized. It has been reported that bipolar cell activation thresholds are lower for transverse cathodic currents than for equivalent anodic currents (Jensen et al. 2003). Our simulations provide insight into this phenomenon. The passive cable segments used in our preliminary study were electronically short and thus can be treated as approximately isopotential. The extracellular potential, on the other hand, grows more negative in the direction of the cathode, here assumed to be on the vitreous surface. At the proximal tip of a vertically oriented process, the extracellular potential will be most negative and thus closest to the intracellular potential (assumed to be uniform), causing the membrane potential at that point to be depolarized from the resting potential. At the distal tip, on the other hand, the difference between the intracellular and extracellular potential is maximal, and the corresponding membrane potential is hyperpolarized. This example shows how the stimulation protocol must be designed in tandem with information about cellular morphology.