

Predicting Every Spike: A Model for the Responses of Visual Neurons

Justin Keat,¹ Pamela Reinagel,² R. Clay Reid,² and Markus Meister^{1,3}

¹Molecular and Cellular Biology
Harvard University
16 Divinity Avenue
Cambridge, Massachusetts 02138

²Neurobiology
Harvard Medical School
Boston, Massachusetts 02115

Summary

In the early visual system, neuronal responses can be extremely precise. Under a wide range of stimuli, cells in the retina and thalamus fire spikes very reproducibly, often with millisecond precision on subsequent stimulus repeats. Here we develop a mathematical description of the firing process that, given the recent visual input, accurately predicts the timing of individual spikes. The formalism is successful in matching the spike trains from retinal ganglion cells in salamander, rabbit, and cat, as well as from lateral geniculate nucleus neurons in cat. It adapts to many different response types, from very precise to highly variable. The accuracy of the model allows a compact description of how these neurons encode the visual stimulus.

Introduction

A central problem of research in neural coding is the relationship between neural activity in the brain and the behaviorally relevant variables of the external world. In the sensory periphery, those variables are the inputs from sense organs. In the visual system specifically, they are completely defined by the light intensity projected on the retina, as a function of wavelength, space, and time. For a given neuron in the visual system, one would like to have a compact description that predicts the spike train in response to any such visual stimulus. If such a neural code could be found, it would specify what aspects of the world this neuron represents and what it could communicate to other neurons in the circuit. Furthermore, the code would specify what aspects of the spike train are used to represent those features. As a result, one could better understand the neuron's role in the information processing task of the entire circuit.

One obstacle to a simple mathematical description of sensory responses is that neurons tend to produce variable spike trains even to identical repeats of the same stimulus. For any given stimulus, there are many possible responses, and a complete neural code would specify the probability that each of these will occur. Obviously, a probability table relating every possible stimulus to every possible spike train would be somewhat unwieldy, and therefore some simplifying assumptions are in order. A popular assumption is that the

neuron generates spikes independently of each other: at every instant it “decides” whether or not to fire a spike, and the probability of firing varies with time as a function of the stimulus. In this limit, the probabilities of getting various spike trains are completely determined by the time course of the instantaneous firing rate $r(t)$ (Rieke et al., 1997; Meister and Berry, 1999). Experimentally, this can be estimated by accumulating a peristimulus time histogram (PSTH) of spikes from many identical trials. The problem of neural coding is then reduced to predicting the firing rate $r(t)$ as a functional of the sensory stimulus $s(t)$ (Rodieck, 1965; Sakuranaga and Naka, 1985; Victor, 1987).

The assumption that sensory responses are fully described by the time-varying firing rate may well be justified in certain brain areas, for example, deep in the cortex (Shadlen and Newsome, 1994). Neurons there generate highly variable spike trains because they receive large numbers of unsynchronized synaptic inputs and because many of these are not under experimental control. By contrast, neurons in the early visual system—from the retina to the lateral geniculate nucleus (LGN) to area V1—can deliver remarkably reproducible spike trains, whose trial-to-trial variability is clearly lower than predicted from the simple firing rate formalism (Berry et al., 1997; Reich et al., 1997; Berry and Meister, 1998; Kara et al., 2000; Reinagel and Reid, 2000). Often, individual action potentials are reproducibly time locked to the stimulus. Given that the synapses in this pathway leading to visual cortex are especially reliable (Reid and Alonso, 1995; Usrey et al., 1999), individual spikes can, in fact, have a strong effect on postsynaptic neurons. Consequently, one needs a different framework for studying this neural code that allows the prediction of individual spikes and spike patterns with high timing accuracy, but also accounts for the remaining stochastic variability in these responses.

When exposed to a visual stimulus rich in temporal variation—such as natural movies or random flicker—retinal ganglion cells typically respond with brief clusters of spikes separated by longer periods of silence (Berry et al., 1997; Berry and Meister, 1998; Meister and Berry, 1999) (Figure 1A). If the stimulus is repeated, the clusters occur reliably at the same times, sometimes to within 1 ms. Thus, the PSTH accumulated over many trials is a series of sharp peaks of high firing rate separated by intervals where the firing rate is absolutely zero (Figure 1B). Similar sharp firing events in response to temporally rich stimuli have been reported in relay cells of the lateral geniculate nucleus (Dan et al., 1996; Reinagel and Reid, 2000) in the nucleus of the optic tract (Clifford and Ibbotson, 2000), in area V1 of the primate visual cortex (Vinje and Gallant, 2000), and even in cortical area MT (Bair and Koch, 1996; Buracas et al., 1998).

The spike clusters produced by retinal ganglion cells are reproducible not only in their timing but also in the number of spikes they contain: the trial-to-trial variance in the spike count of a particular firing event is often less than one and almost always less than the mean spike count (Berry et al., 1997; Berry and Meister, 1998;

³Correspondence: meister@biosun.harvard.edu

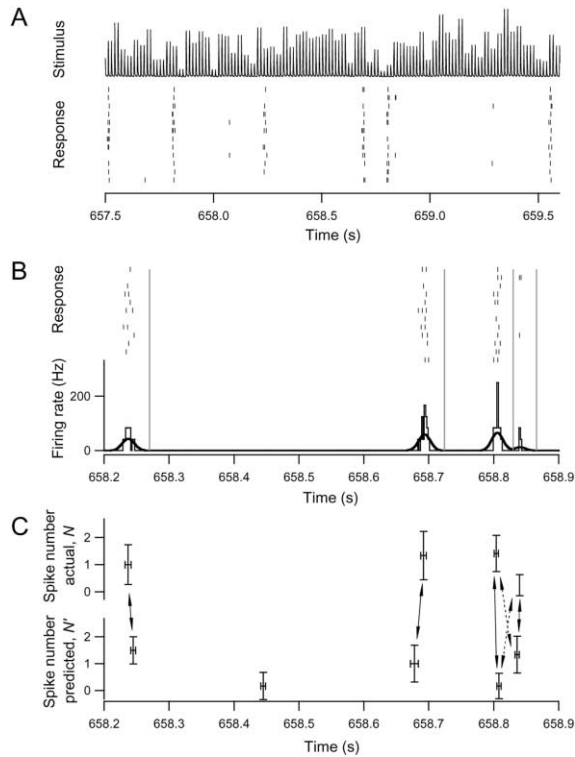


Figure 1. Precise Firing Events in the Response of a Retinal Ganglion Cell

(A) The intensity time course of the flicker stimulus (top trace) and responses of a salamander “strong OFF” cell to 12 repetitions of that stimulus. Each vertical mark represents a spike.

(B) Parsing a set of spike trains into events (see Experimental Procedures for details). The PSTH accumulated across trials (thin curve) is smoothed (thick curve) with a Gaussian filter. When two peaks in the smoothed curve are separated by a significant trough, a division (vertical gray lines) is drawn between the two firing events.

(C) Matching real (top) and predicted (bottom) firing events. Each data point represents event time T and spike number N of the corresponding event. The error bars indicate the variability across trials V of the event time and the variability S of the spike number. Arrows show the correspondences that minimize the error measure. Dotted arrows illustrate an illegal “crossed” correspondence (see Experimental Procedures).

Kara et al., 2000). Thus, spike clusters differing by just one spike can reliably convey different messages about the visual stimulus, and they do so with a timing precision of a few milliseconds. The key variables by which a firing event conveys its visual message are the time of the first spike and the total number of spikes. By comparison, the detailed timing of subsequent spikes within the firing event contributes little to neural coding, carrying only $\sim 5\%$ of the visual information (M. Berry and M.M., unpublished data). However, these prior studies did not address just what features of the stimulus these firing events convey.

Here we develop a mathematical model that can predict neuronal spike trains from the time course of the stimulus. The model is applied to recordings of responses to random flicker stimuli from retinal ganglion cells in salamander, rabbit, and cat, as well as LGN neurons in cat. For each neuron, we optimize the parameters of the model so that the simulated responses

match the real ones. In almost all cases, the algorithm makes successful predictions: the simulated spike trains are about as close to the real spike trains as the real spike trains are across trials. Remarkably, the model can capture the behavior of a wide range of cell types, with seemingly very different light responses.

Results

As discussed above, rapidly varying stimuli elicit precise and reliable responses from many visual neurons. For example, Figure 1 illustrates the response of a salamander retinal ganglion cell to randomly flickering light. The neuron fires at quite precisely defined moments during the flicker sequence (Figure 1A). Each firing event consists of a small cluster of spikes, and subsequent events are clearly separated by periods of complete silence (Figure 1B). Our mathematical model of this process should take the visual stimulus as input and generate spike trains with discrete firing events that closely resemble those of the real neuron. To evaluate this correspondence quantitatively, we measure several salient properties of these firing events in real and simulated spike trains.

Following previous work (Berry et al., 1997; Berry and Meister, 1998), the most important aspects of a firing event are its time of occurrence T , measured as the trial-averaged time of the first spike, and its spike number N , measured as the trial-averaged number of spikes in the event. Though the spike trains appear reproducible, there clearly is some trial-to-trial variability (Figure 1B). Therefore, we also measure the time jitter V as the standard deviation across trials of the time of the first spike, and the number jitter S as the standard deviation across trials of the spike number. In this way, each firing event is characterized by four numbers (T, N, V, S) (Figure 1C; Equation 9 [all equations in Experimental Procedures, below]). The model’s goal is to produce simulated spike trains that match these event properties as closely as possible. To this end, the model’s parameters are adjusted for the optimal fit during a flicker sequence that typically elicits tens to hundreds of firing events. Then the model’s predictions are evaluated on the neuronal responses to a different stimulus sequence. Details of the fitting and testing process are presented in Experimental Procedures. In the following sections, the model is assembled stepwise, and at each stage of improvement, we illustrate its successes and remaining deficits by comparing representative real spike trains with the predictions of the model.

Predicting the Occurrence of an Event

We begin with a simple algorithm to predict the times T when firing events occur, without regard to the number of spikes they contain or the variability across trials (Figure 2A). The stimulus $s(t)$ represents the light intensity as a function of time. This is passed through a linear filter to produce the generator potential $g(t)$. When $g(t)$ crosses a preset threshold from below, an event is fired. When the function passes back through threshold from above, nothing happens.

The only parameters of this model are the filter, given by the impulse response $F(t)$, and the threshold θ [this




Figure 2. Predicting Only the Times of Firing Events
(A) The stimulus $s(t)$ is convolved with a filter $F(t)$ to produce the generator potential $g(t)$. An event is fired when the generator potential crosses upward through the threshold θ .
(B) A brief segment of the fit to responses from a salamander “strong OFF” retinal ganglion cell. $r(t)$ shows the times of the first spike in each event, recorded during a single trial. $r'(t)$ shows the predicted event times when $g(t)$ (solid line) crosses the threshold (dashed line).
(C) A histogram of inter-event intervals for a real cell (thick line) and the model’s fit (thin line), showing that the model predicts too many events that closely follow other events.

version corresponds to $a(t) = b(t) = P(t) = 0$ in Equations 1–4]. Once the filter and threshold are optimized, the correspondence between the real firing events and the predicted ones is for the most part quite strong (Figure 2B). However, the model tends to overpredict the number of events that closely follow a preceding event, such as the last event predicted in Figure 2B. Correspondingly, the number of events separated from the preceding one by a long interval is underpredicted (Figure 2C). It appears that the model could be improved by implementing some partial refractoriness after a firing event.

Predicting All Spikes in an Event

The above scheme can be expanded to generate all of the spikes within firing events by simply adding a negative feedback loop. In the model of Figure 3A, each spike triggers a negative after-potential $P(t)$ that gets added to the generator potential $g(t)$. Thus, the sum $h(t)$ immediately drops below threshold, and if $g(t)$ continues to rise, the model will fire again. Thus, large excursions of $g(t)$ now lead to clusters of several spikes during the rising phase (Figure 3B).

After such a firing event, $h(t)$ is considerably lower than $g(t)$ as a result of the accumulated after-potentials,

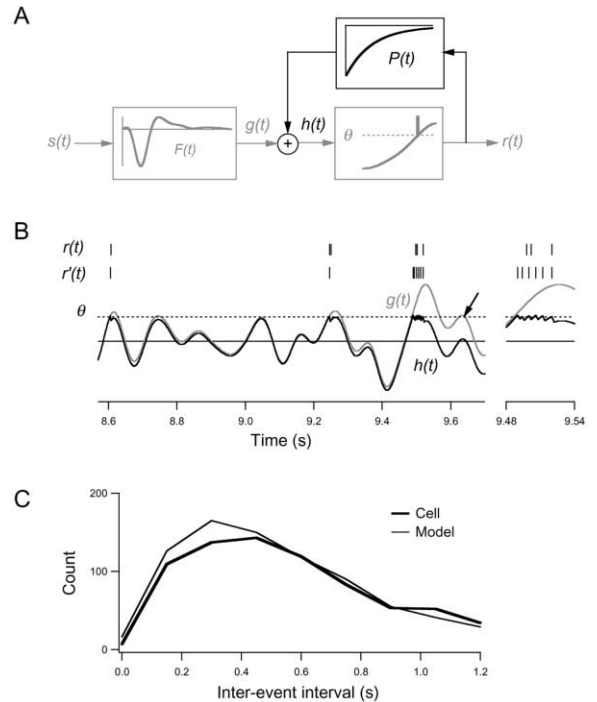


Figure 3. Predicting All Spikes
(A) The spiking model of Figure 2A (gray) augmented by an additional feedback pathway (black). The exponential feedback potential $P(t)$ is triggered by each spike and lowers $h(t)$, the input to the threshold operation.
(B) The fit to responses from a salamander “strong OFF” cell. $r(t)$ shows real spikes from a single trial, and $r'(t)$ the predicted spikes. The expanded time scale on the right shows details of repetitive firing as a result of the feedback potential. The negative feedback also suppresses events that closely follow other events (arrow).
(C) Histogram of inter-event intervals for the real cell (thick line) and the model’s fit (thin line).

and thus the probability for subsequent firing events is temporarily reduced until the after-potentials decay. For example, the false event predicted in Figure 2B no longer occurs in Figure 3B (arrow). The suppression of these events allows the fitting algorithm to lower the model’s spike threshold, which relieves the underprediction of events separated by large intervals. Figure 3C shows that most of the errors in Figure 2C are eliminated by this addition. Thus, the negative feedback mechanism serves both to simulate repetitive firing within a firing event and to implement the refractoriness following an event.

Predicting the Variation across Trials

The last task is to predict the variability of the firing events, specifically the trial-to-trial variations of the event time, V , and of the spike number, S . We did so by adding two Gaussian noise sources, as shown in Figure 4A. Since the noise signals are different from trial to trial, the simulated spike trains now vary across trials as well (Figure 4B).

The first noise source, $a(t)$, is added to the generator potential prior to the threshold. This introduces random variation in the time of the threshold crossing at the

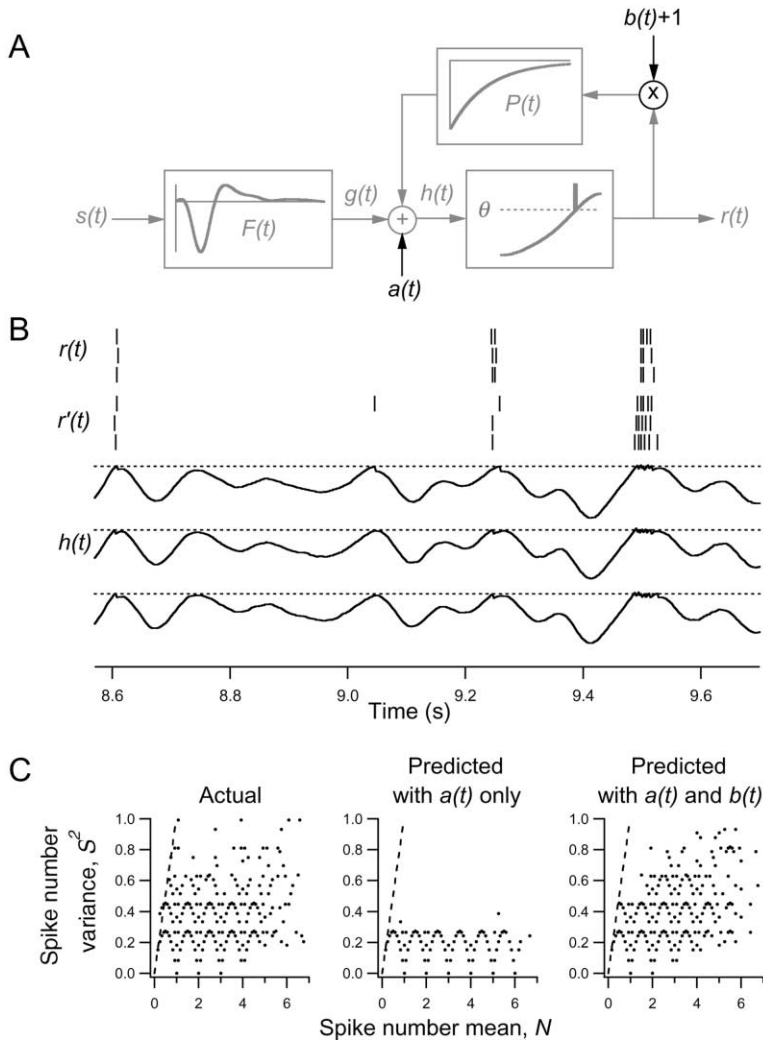


Figure 4. Predicting Variability across Trials (A) The spiking model of Figure 3A (gray) augmented by two noise sources (black). (B) The fit to responses from a salamander “strong OFF” cell. $r(t)$ shows real spikes from three trials, and $r'(t)$ the predicted spikes on three trials. Each simulated trial used different choices for the noise waveforms $a(t)$ and $b(t)$, and thus produced slightly different $h(t)$, the input to the threshold operation. (C) The spike number variance S^2 of firing events plotted against their mean spike number N for a real retinal ganglion cell (left), for the best-fit model using only $a(t)$ (middle), and for the best-fit model using both $a(t)$ and $b(t)$ (right). The arched patterns result from the fact that the spike number is necessarily integer, which constrains the possible values for its variance (Berry et al., 1997). The dashed line corresponds to the identity, the relationship expected for a rate-modulated Poisson spike train.

beginning of a firing event. We chose $a(t)$ to have a Gaussian amplitude distribution with standard deviation σ_a and an exponentially decaying autocorrelation function with time constant τ_a . The time constant was fixed at $\tau_a = 0.2$ s for neurons from salamander and rabbit retina, and $\tau_a = 0.02$ s for cat neurons. These values served very well in predicting the trial-to-trial jitter of the onset of firing events, and holding them fixed helped to reduce the number of free parameters of the model.

In the spirit of maintaining simplicity, we had hoped that this single stochastic component would explain the variability in both event timing and spike number. However, we found that the second and subsequent spikes in a firing event are much more variable than expected from $a(t)$ alone. A model including only $a(t)$ predicts a very low trial-to-trial variation in the spike number of firing events, considerably smaller than that observed in real neurons (Figure 4C). This is because the optimal value for the correlation time τ_a significantly exceeds the duration of most firing events, and thus generation of spikes after the first threshold crossing becomes essentially deterministic. The greater variability of real spike trains could be explained if each spike injects additional noise into the process. We implemented this by invoking a second stochastic component $b(t)$, which

randomly modulates the amplitude of the feedback potential $P(t)$ following each spike. This noise source was taken to have a Gaussian distribution with standard deviation σ_b and a very short correlation time, so its values are independent from spike to spike. With this addition, the model successfully accounts for the spike number variability (Figure 4C).

Simulated Spike Trains

This model was used to fit the responses of neurons in the retina and lateral geniculate nucleus. For each neuron, we identified the physiological cell type by traditional criteria (see Experimental Procedures), recorded the response to a random flicker stimulus, and then optimized the parameters in the mechanism of Figure 4A to match those spike trains. To illustrate the performance of this algorithm, we present raster plots of real spike trains on several identical stimulus trials, along with the corresponding predicted spike trains from the best-fit model (Figure 5). These brief episodes of the response were chosen to have prediction errors typical for their respective cell type. In addition, an analysis of all the firing events produced by one cell is given in Figures 6 and 7, and Table 1 gives summary measures of the model’s performance across all the neurons we

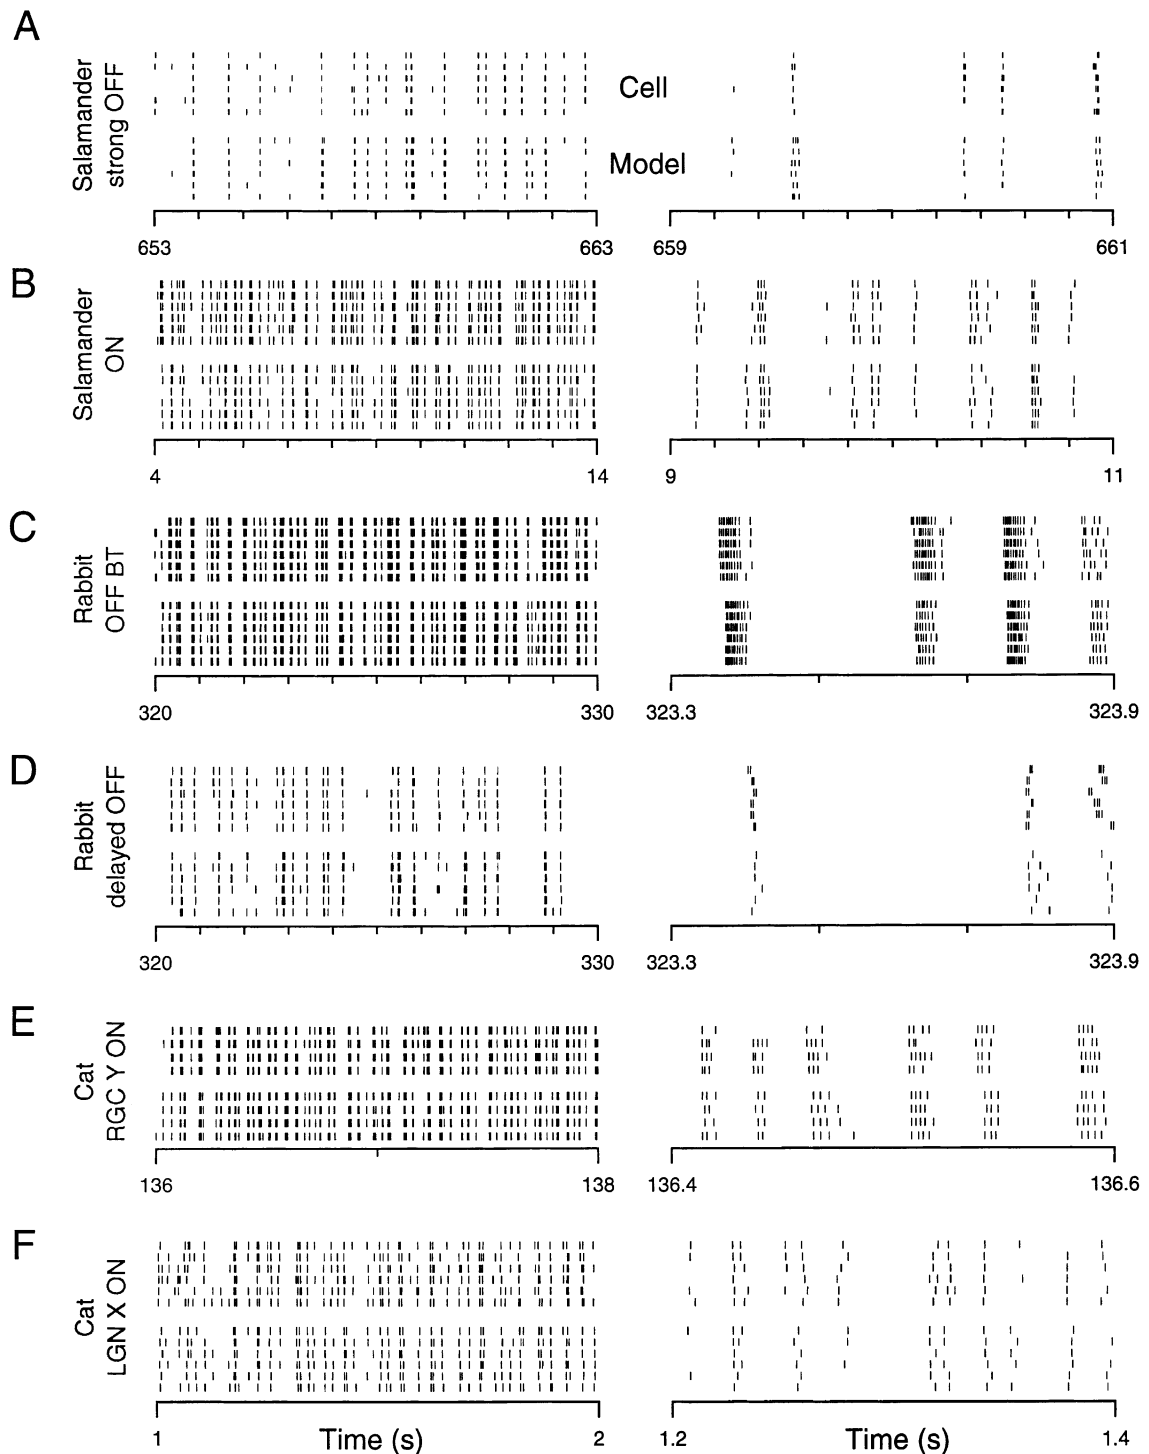


Figure 5. Comparison of Real and Simulated Spike Trains

Each raster plot is for a single cell and shows neural spike trains from several stimulus trials (top half) and corresponding spike trains predicted by the model (bottom half). Graphs on the right show detail on an expanded time axis. The sample cells are from the following types: salamander retina strong OFF (A), salamander retina ON (B), rabbit retina OFF brisk transient (C), rabbit retina delayed OFF (D), cat retina Y ON (E), and cat LGN X ON (F). Note the different time scales: tick intervals are 1 s in left hand panels, 0.2 s in right hand panels.

analyzed. In interpreting these results, recall that the model parameters were always derived from a different stimulus segment than the one used for evaluating the fits; in this sense, the simulated spike trains are truly predictions.

Figure 5A shows responses from a salamander “strong OFF” ganglion cell. These neurons produce very sparse spike trains: the firing rate, averaged over all stimulus repeats, is exactly zero more than 94% of the time because the spikes are locked to the stimulus with

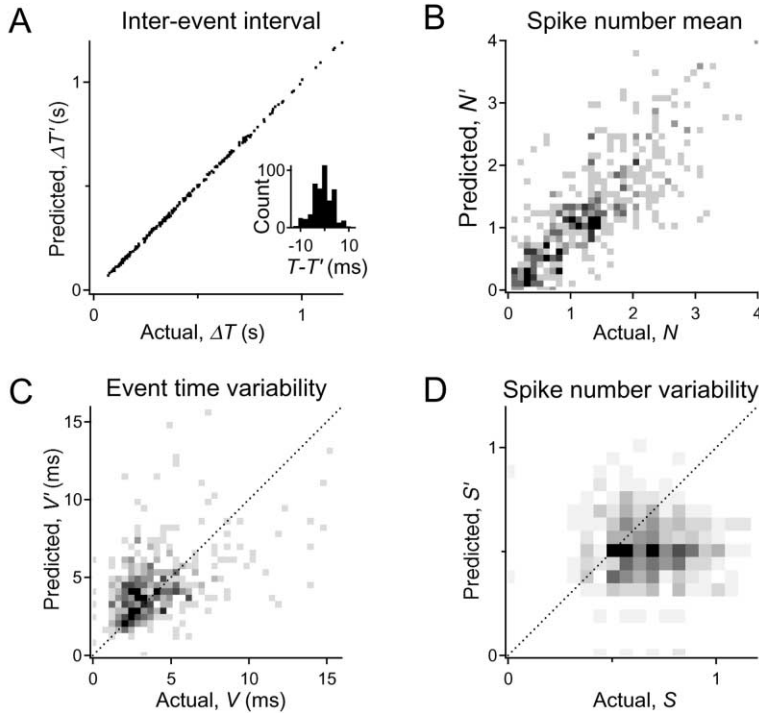


Figure 6. Comparison of Actual and Predicted Firing Events for the Salamander Strong OFF Cell of Figure 5A

(A) Trial-averaged event time; the inter-event interval $\Delta T' = T'_{j+1} - T'_j$ in the predicted event train is plotted against the corresponding interval $\Delta T = T_{j+1} - T_j$ in the actual event train. Inset: histogram of the time difference $T_i - T'_j$ between an actual event i and a matching predicted event j .

(B) Trial-averaged spike number, predicted value N' plotted against actual value N . Panels (B)–(D) are 2-dimensional histograms, with the gray level indicating the number of counts in each bin.

(C) Standard deviation across trials of the event time, predicted value V' plotted against actual value V .

(D) Standard deviation across trials of the spike number, predicted value S' plotted against actual value S .

high timing precision (3.5 ms; see Table 1). The model matches this behavior very closely and predicts correctly the times and spike numbers of almost all firing events. On a finer time scale, Figure 5A shows that the duration of predicted events (the time from first to last spike of an event) is somewhat longer than that of the actual events. Note that the event duration was not a criterion in optimizing the model parameters, but if one includes a term for duration in the error measure, Equation 10, this kind of discrepancy can largely be elimi-

nated (data not shown). A salamander ON cell is illustrated in Figure 5B. These cells fire much less sparsely than the strong OFF cells (Berry et al., 1997): their firing events are more frequent, last longer, and vary more in their timing. Still, the model adjusts to these spiking statistics and produces spike trains that match the real ones both qualitatively and quantitatively (Table 1).

An OFF brisk transient cell from the rabbit retina (Figure 5C) had very different light responses. One finds clear firing events, whose onset is very precise (3.6 ms;




Figure 7. Systematic Error in the Predictions of the Model Compared to the Variability of Neuronal Responses

(A) Error in predicting the mean event time. Each data point is for one cell, with cell types identified by different markers. The discrepancy between the mean time of a predicted event and that of the corresponding actual event was averaged across all the matched events from that cell, $|\overline{T} - \overline{T}'|$. This is plotted against the trial-to-trial jitter \overline{V} of the actual event time, also averaged over all events. Dotted line represents equality.

(B) Error in predicting the mean spike number of an event. As in (A), but plotting the discrepancy $|\overline{N} - \overline{N}'|$ between model and neuron against the trial-to-trial number jitter \overline{S} of the neuron.

(C) Error in predicting the jitter of the event time. As in (A), but plotting the discrepancy $|\overline{V} - \overline{V}'|$ between model and neuron against the neuron's trial-to-trial timing jitter \overline{V} .

(D) Error in predicting the jitter of the spike number. As in (A), but plotting the discrepancy $|\overline{S} - \overline{S}'|$ between model and neuron against the neuron's trial-to-trial number jitter \overline{S} .

Table 1. Error Measures for the Model's Performance, Averaged over All Neurons of a Given Cell Type

System	Cell Type	Number of Cells	\bar{V} (ms)	\bar{S} (spikes)	$ \bar{T} - \bar{T}' $ (ms)	$ \bar{N} - \bar{N}' $ (spikes)	$ \bar{V} - \bar{V}' $ (ms)	$ \bar{S} - \bar{S}' $ (spikes)
Salamander retina	ON	6	18.0	0.41	17.8	0.47	8.23	0.17
	strong OFF	15	3.50	0.58	3.39	0.56	1.45	0.21
Rabbit retina	ON BT	1	5.95	0.93	10.7	1.47	3.20	0.29
	OFF BT	5	3.63	0.87	6.13	2.17	1.64	0.27
	sustained OFF	1	5.88	0.51	8.12	0.90	2.28	0.24
	delayed OFF	2	6.33	0.92	6.88	1.17	2.98	0.40
Cat retina	Y ON	1	0.78	0.70	1.16	1.40	0.45	0.43
Cat LGN	X ON	3	1.58	0.45	1.55	0.36	0.93	0.17
	Y OFF	1	1.03	0.47	1.07	0.54	0.76	0.24

\bar{V} and \bar{S} denote the neuron's variability in event timing and spike number, respectively, averaged over all firing events. These set the scale for precision in the real visual response, on which the accuracy of the model's prediction should be evaluated. Listed to the right are the discrepancies between the predicted and actual event trains in terms of timing $|\bar{T} - \bar{T}'|$, spike number $|\bar{N} - \bar{N}'|$, timing variability $|\bar{V} - \bar{V}'|$, and number variability $|\bar{S} - \bar{S}'|$. In each case, the absolute value of the discrepancy was averaged over all events. Note these error measures are those used in optimizing the model (Equation 11).

see Table 1), but they contain large numbers of spikes (often >10), with low trial-to-trial variability (<1 spike; see Table 1). Again, the model captures the timing, spike number, and overall shape of these firing events quite accurately. By contrast, a delayed OFF cell (Figure 5D) responded very sparsely, with a zero firing rate more than 91% of the time. Note that its firing events align with a subset of the events of the OFF brisk transient cell (Figure 5C). The event timing, spike number, and the variability of those quantities are again predicted rather well by the model. However, the expanded time scale in Figure 5D reveals an additional form of trial-to-trial variability in these spike trains: a brief cluster with very constant spike number jitters back and forth in time by an amount greater than that of the cluster duration. The firing mechanism of Figure 4A cannot produce this type of behavior for the following reason: the time of the first spike in a cluster can jitter from trial to trial depending on the value of the noise $a(t)$. However, $a(t)$ varies slowly and remains essentially constant for the remainder of the event. Therefore, the model will continue firing spikes until the generator potential reaches its peak (see the large cluster in Figure 3B). Thus, the spike cluster should terminate at approximately the same time in each trial, counter to what happens in Figure 5D. In principle, one could capture this form of variability by allowing the latency of the filter to fluctuate, which would produce the same spike cluster with slightly different timing on each trial.

A Y-type ON ganglion cell from the cat retina (Figure 5E) produced firing events at a much higher rate than that of the neurons considered so far and with great timing precision (0.8 ms; Table 1). The model adapts to these dynamics, and it is difficult to distinguish the set of simulated spike trains from the real ones. An X-type ON cell from the cat LGN (Figure 5F) fired events at an even greater frequency. However, each event contained fewer spikes on average, and many events were represented on only a subset of the trials. Again, most aspects of this cell's behavior are well matched by the predicted spike trains.

Evaluation of Performance

A more quantitative evaluation of the simulated spike trains is given in Figure 6 by comparing the actual and

predicted event trains for an individual cell, and a summary over many neurons is shown in Figure 7. The timing of firing events is reproduced very well by the model (Figure 6A): in comparing the inter-event intervals in actual and simulated spike trains, the discrepancies appear negligible relative to the inter-event intervals themselves. In fact, the root-mean-square timing error incurred by the model is only 3.5 ms (Figure 6A inset). This is comparable to the amount by which this cell's event time jittered between trials, 3.9 ms. A similar level of accuracy was achieved for many cells (Figure 7A; Table 1). In almost all cases, the model predicted the timing of firing events with errors comparable to the trial-to-trial jitter. Exceptions are the brisk transient cells in the rabbit retina, where the model's errors are up to twice as large as the timing variability of the cell. Remarkably, the firing times of Y cells in both cat retina and LGN were matched correctly to within 1.1 ms.

The number of spikes in each event was also predicted very accurately (Figure 6B), with the discrepancy between data and prediction often less than one spike. For most cells, the error in predicting the spike number was comparable to the neuron's variation across trials (Figure 7B; Table 1). Again, the predictions for rabbit brisk transient cells fall short of this mark. Note, however, that these firing events contain large numbers of spikes (e.g., Figure 5C) so that the errors of the prediction still constitute a small fraction of the absolute spike numbers.

The model also served to predict the trial-to-trial jitter of the event time V and of the spike number S . For the neuron in Figure 6, both the average time jitter and the average number jitter are matched correctly. Event by event, there is a positive correlation between the actual V and the predicted V' (Figure 6C), but little between the actual S and predicted S' (Figure 6D). However, the number jitter S simply does not vary much across events; what variation there is depends largely on how close the average number of spikes N is to an integer, with S smaller for near-integral values of N (Figure 4C). So to match S more accurately, the model would need to predict N to within a small fraction of a spike. Nevertheless, across all the cells in this study (Figures 7C and 7D; Table 1), the model reproduced the stochastic fluctuations V and S correctly to within $\sim 50\%$.

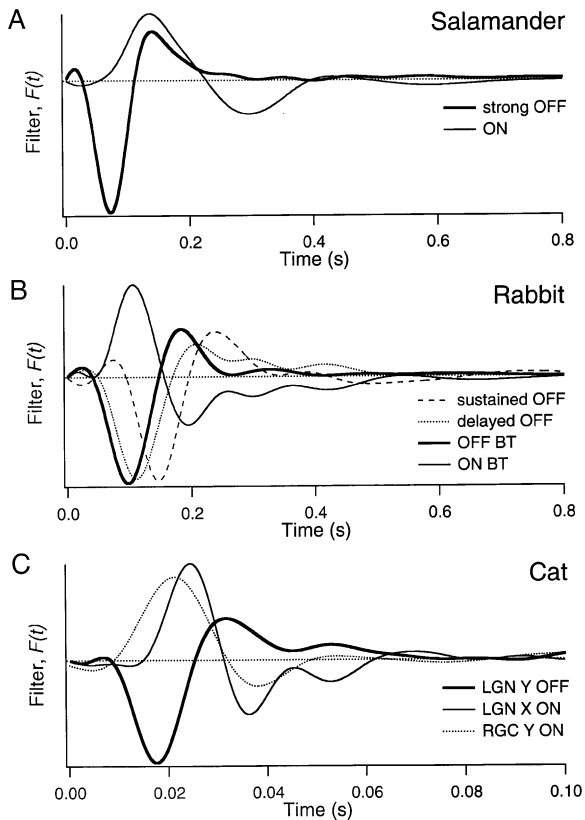


Figure 8. Comparison of Filter Functions across Cell Types
The waveform of the model's filter $F(t)$, averaged across all cells of the same type from salamander (A), rabbit (B), and cat (C).

Discussion

What is the purpose of constructing a mathematical model for neuronal responses? On a practical level, the model allows for a concise description of individual neurons. Each cell's visual response is characterized by a small set of numbers, the model parameters (Figure 8; Table 2). These numbers are communicated easily to another researcher, who may use them to implement a realistic neuron of that type. For example, simulated retinal spike trains can serve as input to computational models of the visual system. In the distant future, the simulation might even be performed by a neural prosthesis to replace retinal circuitry and drive optic nerve fibers directly (Humayun et al., 1999).

On a deeper level, an accurate model for neural responses also serves as a statement of the neural code. It specifies the statistical relationships between stimuli and responses. Given a successful "forward" description of firing as a function of the stimulus, one can derive the corresponding "reverse" description: for any given spike train, what is the probable time course of the light intensity, as well as the uncertainty in that estimate? Such a reverse dictionary would specify what visual features are encoded by the neuron and transmitted to its postsynaptic targets.

Strictly speaking, one can be sure of this code only within the particular visual ensemble in which it was

measured. In exploring whether an accurate prediction of spike trains is feasible, we used a visual environment rich in temporal structure but with no spatial variation. The conclusions in the following sections are subject to this caveat. Extensions of the approach to more general conditions will be discussed.

Quality of the Model's Predictions

In evaluating the accuracy of the model's performance, we placed particular value on the prediction of two aspects of the visual response: the time of occurrence of each firing event and the number of spikes produced in that event. This choice is not tied to any particular theory of neural signaling; most would probably agree that it is important when a cell fires and also how many spikes it fires. By these criteria, the formalism performs very well. In most cases, the error the model makes in predicting the average event time is comparable to the random variation of that event time across trials (Figure 7; Table 1). Similarly, the error in predicting the average spike number compares well to the jitter of that spike number across trials. In other words, the prediction one gets from this model about the timing and spike number of a cell's firing events is as reliable as direct observation of that same neuron on a previous trial with the same stimulus.

Is this level of accuracy sufficient? In the real world, the visual system operates exclusively on single trials, without the luxury of improving resolution by averaging many responses to identical stimuli. Nor is there much opportunity to average across equivalent cells, because neurons in the early visual system tend to tile the visual field with little redundancy (Wässle and Boycott, 1991; Meister and Berry, 1999). Consequently, operation of the visual system under natural conditions does not require the properties of these neurons to be specified more precisely than their trial-to-trial fluctuations. To understand a neuron's role in visual behavior, we therefore suggest that a model of the light response can be deemed successful if its systematic errors are as small as the neuron's random errors.

A Simple Neural Code

Given the accuracy of the model, one can understand how these early visual neurons encode the stimulus, simply by inspecting how the model does it. The qualitative nature of that code seems rather simple (Figure 9). The neuron is "interested" in just one aspect of the visual input: at any point in time, the cell only asks to what extent the recent stimulus looks like the filter $F(t)$. Thus, $F(t)$ identifies *what* feature the cell reports. When the strength $g(t)$ of that feature rises above threshold, the cell fires. So the timing of the first spike in a firing event signals *when* that feature happened. The number of spikes in the subsequent burst depends on the peak excursion of $g(t)$. Therefore, the spike number in the event specifies *how much* of the feature was present in the stimulus. This simple interpretation holds strictly when the neuron's firing events are separated sufficiently in time, by several decay times of the after-potential. At shorter delays from the preceding event, the accumulated after-potential effectively raises the threshold for feature detection.

Table 2. The Best-Fit Parameters of the Model, Averaged over Neurons of the Same Cell Type and Quoted as Mean \pm Standard Deviation

System	Cell type	$F(t)$	θ	τ_p (ms)	B	σ_a	σ_b	τ_a (ms)
salamander retina	ON		0.59 \pm 0.37	110 \pm 22	1.51 \pm 0.73	0.30 \pm 0.16	0.085 \pm 0.084	200
	strong OFF		1.46 \pm 0.27	211 \pm 46	0.72 \pm 0.32	0.15 \pm 0.06	0.33 \pm 0.09	200
rabbit retina	ON BT		-0.13	130	0.74	0.16	0.52	200
	OFF BT		0.11 \pm 0.32	166 \pm 28	0.17 \pm 0.03	0.075 \pm 0.038	0.33 \pm 0.39	200
	sustained OFF		0.53	132	0.93	0.19	0.43	200
	delayed OFF		0.96 \pm 0.14	172 \pm 66	0.47 \pm 0.36	0.17 \pm 0.06	0.57 \pm 0.77	200
cat retina	Y ON		0.10	20	0.65	0.14	0.072	20
cat LGN	X ON		0.76 \pm 0.73	24 \pm 12	4.59 \pm 2.00	0.44 \pm 0.28	0.15 \pm 0.15	20
	Y OFF		1.49	21	4.57	0.35	0.08	20

The filter $F(t)$ (displayed in more detail in Figure 8); the threshold θ ; the decay time of the after-potential τ_p ; the amplitude of the after-potential B ; the noise affecting the generator potential σ_a ; and the noise affecting the after-potential σ_b . The time constant of the noise, τ_a , was not optimized for each individual cell, but is stated here for completeness. Note that θ , B , and σ_a are all expressed relative to the root-mean-square value of the generator potential $g(t)$, and σ_b is expressed as a fraction of B (see Figure 4 and Experimental Procedures).

Remarkably, this description accounts for the responses of many types of neurons, even though their spike trains at face value have very different appearance (Figure 5). Thus, the seeming complexity of response types may reduce to quantitative variants of a single common description, each identified by a specific set of model parameters. Table 2 and Figure 8 present the average parameters obtained for neurons in different cell classes. Because the simulation algorithm is relatively simple, one can understand the role that each parameter plays in creating the characteristic spiking patterns (Figure 5), and we briefly illustrate this correspondence.

The filter $F(t)$ determines what temporal features in

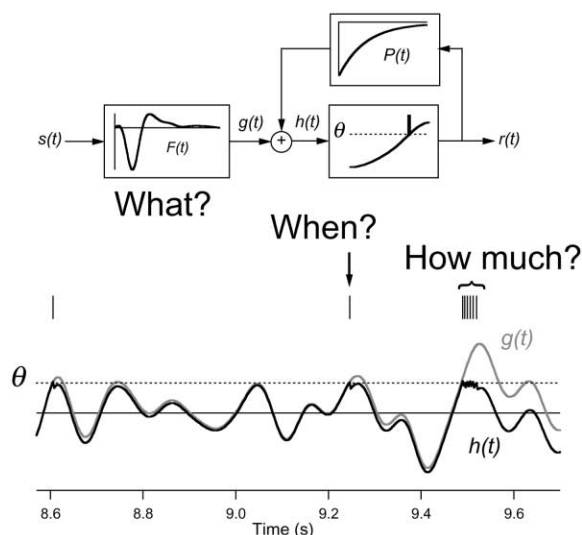


Figure 9. Summary of the Neural Code Implemented by the Model. The filter $F(t)$ encompasses what stimulus feature the spikes represent; the precise onset of a firing event encodes when that feature happened; and the number of spikes in the subsequent cluster reports how strong the feature was.

the flickering stimulus are effective in driving the neuron; a stimulus waveform that matches the time-reverse of the filter will produce a large output more likely to cross threshold. The filters derived for the various cell types (Figure 8) are all biphasic curves, with a strong primary peak followed by a second peak of opposite sign. Thus, all these visual neurons are more sensitive to a change in the light intensity than to a steady maintained level. The curves also differ in many respects: the sign of the principal peak determines whether the response is ON-type or OFF-type; the ratio of the two peaks influences how transient the response is; and the overall time scale of the curve sets the time scale of the response. The filters of rabbit and salamander ganglion cells are comparable in width, whereas those of neurons in cat retina and LGN are considerably faster. The same progression of time scales is evident in the raster plots of Figure 5.

Whereas $F(t)$ sets the overall time scale of the fluctuations in the generator potential $g(t)$, only some fraction of those transients will cross the threshold. Thus, the threshold θ determines the sparseness of firing events in the neuron's response. Once the threshold is crossed, the after-potential immediately reduces the input to the spike generator by an amount B . If the generator potential continues to rise by more than this, subsequent spikes can be fired. Thus, the parameter B , along with the threshold θ , controls the number of spikes fired in an event. For example, the cat LGN cells, which tended to fire in single spikes (Figure 5F), had by far the largest values of B , whereas the rabbit OFF brisk transient cells (Figure 5C), with many spikes per event, had the smallest values (Table 2). The after-potential then decays with time constant τ_p . Because τ_p is usually long compared to the duration of a firing event, this decay does not influence the dynamics of firing during an event. Instead, it determines the duration of partial refractoriness after a firing event. For example, salamander OFF cells have a large τ_p , which contributes to making their firing events sparser than those of ON cells (Figures 5A and 5B).

The trial-to-trial variability of event timing is determined by σ_a , the noise component of the generator po-

tential, and the shape of the filter $F(t)$, which sets the overall time scale of the firing process. For example, salamander OFF cells have a faster filter and a smaller noise component σ_a than those of ON cells (Table 2) and correspondingly much more reliable spike timing (\bar{V} in Table 1). Cat neurons have a fast filter, and consequently low timing jitter, even though σ_a is comparable to that of salamander neurons. The variability of the spike number in an event is determined by both σ_a and the amount of noise that is added with the after-potential following each spike, $B \cdot \sigma_b$. Across cell types, the ratio of the two terms varies a great deal, but the second term often dominates over the first. In practice, therefore, σ_a determines timing precision and σ_b most of the number precision.

Interpretation of the Model

On a formal level, the spike train simulator of Figure 4 is related to the popular “leaky integrate-and-fire” mechanism (Knight, 1972; Fohlmeister et al., 1977; Lankheet et al., 1989; Reich et al., 1998) as follows. One can view $h(t)$ as the membrane potential of a neuron with capacitance C and leak resistance $R = \tau_p/C$. The capacitance gets charged by a synaptic current, obtained by passing the stimulus $s(t)$ through a linear filter. When $h(t)$ crosses the firing threshold, the cell fires a spike. As a byproduct of the spike, the membrane is discharged briefly. This reduces the membrane potential by an amount B , which subsequently decays with time constant $\tau_p = RC$. In this picture, the noise source $a(t)$ represents noise in the synaptic input, and $b(t)$ is noise in the membrane machinery that affects the degree of discharge. Reich and colleagues (Reich et al., 1997; Reich et al., 1998) analyzed the firing statistics of such a noisy leaky integrate-and-fire model and concluded that it was consistent with the observed spiking statistics of cat retinal ganglion cell responses. Here we have shown that by providing the appropriate input to such a spike generator, one can in fact predict the neuron’s entire spike train from the stimulus.

In this picture, the model parameters acquire specific biophysical meaning. The linear filter comes to represent all of retinal processing and dendritic integration involved in producing the cell’s synaptic current. The shape of this function likely reflects the time course of phototransduction in cones, with additional high-pass filtering at subsequent retinal synapses (Warland et al., 1997). The threshold and after-potential components would represent the neuron’s spike-generating machinery (Kistler et al., 1997). In reality, the separation between linear and nonlinear processing is probably not quite as sharp: for example, intracellular recordings from ganglion cells show that the membrane potential is not a linear function of the stimulus (Baylor and Fettiplace, 1979). During stimulation with random flicker, the generator potential already exhibits discrete depolarizing transients, which in turn cause the clusters of spikes in firing events (Sakuranaga et al., 1987). Furthermore, the model is equally successful with LGN neurons as with ganglion cells, even though the former receive their input from the latter. Thus, it is plausible that both the linear and nonlinear stages of the model are somewhat distributed throughout the circuit and do not map uniquely

onto discrete circuit elements. To learn more about this correspondence, it will be useful to record signals from retinal interneurons, such as bipolar and amacrine cells, and compare them to the internal variables of the model.

Choice of Formalism

There have been many attempts to capture the mathematical relationship between a sensory stimulus and neuronal firing. With rare exceptions, they seek to predict the neuron’s average firing rate $\langle r(t) \rangle$, a continuous function of time measured by accumulating the PSTH over many trials. Generally the firing rate is cast as a functional of the stimulus $\langle r(t) \rangle = R[s(t); \alpha_i]$ and then the parameters α_i are optimized to approximate the measured relationship between $\langle r(t) \rangle$ and $s(t)$. Several efficient methods have been developed to find a good functional $R[]$. The most systematic approach writes $R[]$ as a Wiener kernel expansion in the stimulus; its parameters are the Wiener kernels of the response, which can be computed by correlating the response with the stimulus (Marmarelis and Marmarelis, 1978; Victor and Shapley, 1980; Sakai et al., 1988). Another method—less general, but often more efficient—writes $R[]$ as a sequence of elementary operations performed on the stimulus. Each stage in such a “cascade model” performs a simple task on its input, such as temporal filtering or spatial pooling, or a nonlinear memory-less transform (Rodieck, 1965; Victor and Shapley, 1979; Hunter and Korenberg, 1986; Sakai and Naka, 1987; Victor, 1987; Sakai and Naka, 1995; Benardete and Kaplan, 1997; reviewed in Meister and Berry, 1999).

The present approach is somewhat different. The model implements an explicit point process that predicts a spike train with specific spike times. We chose this form for two reasons:

(1) The average firing rate $\langle r(t) \rangle$ is not really a continuous function. These visual neurons—at least when driven with rich stimuli such as random flicker or natural scenes—fire in discrete clusters that are timed reproducibly from trial to trial. Consequently, the average firing rate $\langle r(t) \rangle$ is exactly zero much of the time but rises and falls sharply within a few milliseconds at well-defined instants (Figures 1 and 5) (Berry et al., 1997). It is difficult to model such a series of sharp transients with a formalism designed for continuous functions. By contrast, the formalism of Figure 4A produces such firing events naturally. In the limit of very low noise, spikes will happen at precisely the same time on every trial; with very high noise, the spike times will vary a great deal, and one obtains a smoothly varying firing rate.

(2) It is not sufficient to predict the average firing rate $\langle r(t) \rangle$. To understand what a neuron encodes during sensory processing, one needs to know not only its average response but also how much the response can vary about the average (Meister and Berry, 1999). Early visual neurons can be remarkably precise (Berry et al., 1997; Reich et al., 1998; Kara et al., 2000), but the amount of trial-to-trial variability still depends greatly on the cell type and on the visual stimulus (Figure 7). Thus, a satisfying model of the light response needs to predict not only the mean firing pattern, but also its stochastic properties. An efficient way of doing that is to model the stochastic point process itself. By simulating many trials

and analyzing them in the same way as real neuronal spike trains, one can predict the average firing rate, but in addition many other statistics of the response.

Optimizing the Fit

An important component of this modeling strategy is the method for optimizing the parameters, in particular the choice of the error measure that gets minimized during fitting. This expression E (Equations 10–14) evaluates the discrepancy between the actual set of spike trains and a corresponding predicted set. Its virtue is that it explicitly contends with the discrete firing events in the responses. To appreciate this, it helps to consider a commonly used alternative: in many studies that seek to predict the time course of a neuron's firing rate, the error is measured as the mean squared difference between the actual and the predicted rate (Marmarelis and Marmarelis, 1978; Victor and Shapley, 1980; Sakuranaga et al., 1987; Victor, 1987; Korenberg et al., 1988; French and Korenberg, 1989, 1991; Kondoh et al., 1991; Sakai and Naka, 1995; Benardete and Kaplan, 1997). Suppose now that the actual firing rate has a sharp peak, corresponding to a 10 ms wide firing event (e.g., Figure 1B), and the model with a certain parameter set is able to predict a similar peak but displaced by 10 ms. Since the two firing rate functions have no overlap, the mean squared error is maximal; in fact, the error would be smaller if the model had not predicted any firing at all. Even though the 10 ms timing error could be reduced by a small change in the parameters, it is difficult to reach that optimum by minimizing the mean squared error. By contrast, our error measure E would recognize the correspondence between the actual and predicted event even if they do not overlap in time. The error then decreases smoothly with the timing difference, and thus the model parameters can be optimized successfully.

A crucial step in evaluating this error measure is to decide how events in the actual response should be matched to events in the simulation. The algorithm for finding this correspondence, and indeed the entire form of E , was heavily inspired by the “distance metric” of Victor and Purpura (1996). This is a measure to tell how different two individual spike trains are. It relies on the simple exercise of turning one spike train into the other by sliding spikes in time and creating or destroying spikes. A cost is assigned to each of these elementary operations, and one can find the transformation with the lowest overall cost, which in turn is taken as the distance between the two spike trains. Our distance function E can be seen as an extension of this approach from single spikes to firing events. Its flexible form (Equation 10) allows the model to fit not only the timing of firing events, but many other properties as well.

Regime of Validity

It will be important to test how far this description of visual signaling extends. To limit the computational effort, the present study was restricted to temporal variation in the stimulus. The model can certainly be extended to include spatial processing, for example, by making the first element a space-time filter $F(x,t)$. This will allow the prediction of responses to stimuli $s(x,t)$ that vary in space as well as time. It will also be instructive to explore

very different stimulus environments, including those that do not produce precise firing events. For example, when exposed to constant light, retinal ganglion cells tend to fire at a maintained rate with variable interspike intervals. If a slow sinusoidal variation is superposed on the background, the firing rate is modulated smoothly around the mean (e.g., Enroth-Cugell and Robson, 1966). The model of Figure 4A can certainly replicate these modes of activity. For example, a negative threshold leads to maintained firing, and a high noise level will produce variable intervals. However, it is likely that the same neuron will be described by different sets of parameters under different conditions of stimulation. It will be instructive to explore how the parameters are modulated by recent visual experience, for example, following changes in the mean light level or in the average contrast. The compact parametrization of light responses afforded by the present model promises an equally concise summary of the effects of adaptation. If this can be achieved, one might obtain an accurate model that generalizes even to the stimulus ensemble of ultimate interest: real natural vision.

Experimental Procedures

Overview

We presented a random flicker stimulus to the visual system, repeated identically several times, while recording the spike trains of one or more visual neurons. A mathematical model with 20 parameters was developed to predict spike trains from the stimulus. The parameters were adjusted to optimize the fit to the real spike trains over one half of the experiment. Then the quality of the model was evaluated by how well it predicted responses to different stimuli in the other half.

Recording and Stimulation

Salamander and Rabbit Retina

We recorded extracellularly from retinal ganglion cells of larval tiger salamanders and rabbits as described (Meister et al., 1994; Smirnakis et al., 1997). Briefly, the dark-adapted retina was isolated under dim red light. A piece of retina ~ 2 mm (salamander) or ~ 4 mm (rabbit) in diameter was placed ganglion cell side down on an array of 61 electrodes in a bath of oxygenated Ringer's (salamander) or Ames (rabbit) medium. The electrodes recorded action potentials from nearby ganglion cells. On occasion, more than one neuron contributed to the signal on a given electrode; their spike waveforms were sorted by their shape to obtain reliable single-neuron spike trains.

The stimulus consisted of spatially uniform white light from a computer monitor. A new intensity was chosen at random every 30 ms from a Gaussian distribution with a standard deviation equal to 35% of the mean level. The mean light intensity was ~ 4 mW/m² at the retina, in the regime of photopic vision. Six different stimulus segments were used, either 200 s (salamander) or 80 s (rabbit) in duration. Each of these segments was repeated 12 times, followed by the next segment. The actual time course of the light emitted by the monitor was measured using a photocell (Figure 1A), and this stimulus was used as input for the modeling process.

Cat Retina and LGN

For recordings from cat neurons, the animal was anesthetized (ketamine HCl 20 mg/kg IM, followed by sodium pentothal 20 mg/kg IV supplemented as needed), ventilated through an endotracheal tube, and paralyzed with Norcuron (0.3 mg/kg/hr IV). EKG, EEG, temperature, and expired CO₂ were monitored continuously. Eyes were refracted, fitted with appropriate contact lenses, and focused on a tangent screen. For LGN experiments, we performed a 5 mm diameter craniotomy and recorded single neurons in the A laminae with plastic-coated tungsten electrodes (AM Systems, Everett, WA). For retinal ganglion cell recordings, the electrode was advanced through a guide tube that penetrated the sclera. Voltage signals were ampli-

fied, filtered, and digitized (DataWave Discovery software, Longmont, CO). Spikes from single units were sorted by their waveforms and spike times determined to 0.1 ms resolution.

The stimulus again consisted of spatially uniform white light from a computer monitor, at a mean luminance of 35 cd/m² (LGN) or 50 cd/m² (retina), corresponding to retinal intensities of ~15 mW/m² and ~21 mW/m², respectively. A new intensity value was chosen at random every 7.8 ms, either from a Gaussian distribution with standard deviation equal to 49% of the mean (LGN) or from a binary distribution with 100% contrast (retina). Stimulus segments lasted 63.6 s with 6–8 repeats (LGN) or 31.8 s with 4 repeats (retina). The recorded spike trains did not lock to the video frame rate (128 Hz), as judged from inspecting their power spectra.

All surgical and experimental procedures were in accordance with NIH and USDA guidelines and were approved by the Harvard Medical Area Standing Committee on Animals.

Cell Classifications

Salamander retinal ganglion cells were classified into functional types by the time course of their reverse-correlation to the flicker stimulus (Warland et al., 1997), in particular the time-to-peak. In seven preparations, we distinguished the following types: strong OFF (time-to-peak ≈ 65 ms), weak OFF (82 ms), and ON (107 ms). The weak OFF cells generally had firing rates too low for reliable analysis.

Rabbit retinal ganglion cells were classified by the criteria of DeVries (DeVries and Baylor, 1997; DeVries, 1999), using the measured auto-correlation and reverse-correlation functions, and responses to uniform light flashes. Since our conditions of light adaptation were different, these identifications should be regarded as tentative. In three preparations, we encountered the following types: ON brisk transient, OFF brisk transient, OFF sustained, OFF delayed, ON-OFF direction selective.

Neurons in the cat retina and LGN were classified as X- or Y-type based on their spatiotemporal receptive fields and responses to contrast-reversing gratings. Putative Y cells had large receptive field centers and short latencies relative to other neurons at the same eccentricity. The retinal ganglion cells were also subjected to the null test of Hochstein and Shapley (1976).

Model to Predict Responses

Here we describe the final form of the model, with attention to mathematical and computational details. The complete algorithm is shown schematically in Figure 4A. Formally, the predicted firing rate of the neuron is given by

$$r(t) = \delta(h(t) - \theta) \dot{h}(t) H(\dot{h}(t)), \quad (1)$$

where

$$h(t) = g(t) + a(t) + \int_{-\infty}^t r(\tau)(1 + b(\tau)) P(t - \tau) d\tau \quad (2)$$

$$g(t) = \int_{-\infty}^t s(\tau) F(t - \tau) d\tau \quad (3)$$

$$H(x) = \begin{cases} 1, & \text{if } x > 0 \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

The function $r(t)$ is a series $\sum_i \delta(t - t_i)$ of delta function spikes that happen at times t_i when the generator potential $h(t)$ crosses the threshold θ in the upward direction. $h(t)$ in turn has several components: $g(t)$ is a filtered version of the stimulus time course $s(t)$, obtained by convolution with the filter function $F(t)$. $a(t)$ is a Gaussian noise source. The final term contains a transient feedback triggered by each spike, a convolution of the spike process $r(t)$ with the feedback potential $P(t)$, modulated in amplitude by another Gaussian noise source $b(t)$.

The filter function $F(t)$ defines to a large extent when firing events will occur, and most of the model's parameters are dedicated to adjusting this function. To capture the time course efficiently, we expanded the filter in an orthonormal basis set,

$$F(t) = \sum_{j=1}^N k_j f_j(t), \quad (5)$$

where the k_j are the free parameters. In choosing the basis functions $f_j(t)$, the Fourier basis set was found unsatisfactory, requiring as

many as 28 parameters to produce filters with the desired accuracy. A more efficient choice was made by considering the properties of the typical filter function (see Figure 8): at short times it has high amplitude and varies rapidly, but at longer times it tails off gently. Thus, we used a set of sine functions that were stretched at long times:

$$f_j(t) = \begin{cases} \sin\left(\pi j \left(2 \frac{t}{\tau_F} - \left(\frac{t}{\tau_F}\right)^2\right)\right), & \text{for } 0 \leq t \leq \tau_F. \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

The maximal length of the filter was $\tau_F = 0.95$ s (salamander and rabbit), $\tau_F = 0.25$ s (cat retina), or $\tau_F = 0.19$ s (cat LGN). These functions were then orthonormalized using the Gram-Schmidt method. In contrast with the Fourier basis, at most 16 of these components were required for an accurate fit of the filters.

The shape of the after-potential $P(t)$, which lowers $h(t)$ after every spike and creates a relative refractory period, was chosen as

$$P(t) = B \exp(-t/\tau_p), \quad (7)$$

where B and τ_p are free parameters.

The noise signals $a(t)$ and $b(t)$ both have Gaussian amplitude distribution with zero mean and standard deviations σ_a and σ_b , respectively. $a(t)$ has an exponential auto-correlation function with a time constant of $\tau_a = 0.2$ s (salamander and rabbit) or $\tau_a = 0.02$ s (cat). $b(t)$ is evaluated only at spike times (see Equation 2) and chosen independently for each spike. σ_a and σ_b are free parameters.

A final parameter is the threshold θ . Note that increasing θ is equivalent to reducing the amplitudes of $F(t)$, $P(t)$, and $a(t)$ all by the same factor (see Equations 1 and 2). This redundancy was resolved by normalizing $F(t)$ such that the filtered stimulus $g(t)$ had a standard deviation of 1. Thus, the values of θ , B , and σ_a are all expressed relative to the root-mean-square amplitude of the stimulus-dependent input to the spike generator $g(t)$. In summary, the 20 parameters of the model are used as follows: 15 specify the shape of the filter $F(t)$, one the spike threshold θ , two the after-potential $P(t)$, and two the standard deviations of the noise sources $a(t)$ and $b(t)$.

For numerical evaluation, we used discrete time steps of 2 ms (salamander), 1 ms (rabbit), or 0.2 ms (cat). Simulation on several trials, using the same stimulus but different noise functions, produced a set of spike trains that was compared with the actual spike trains on several stimulus repeats. The gradual optimization of the model parameters was computationally intensive, requiring the above equations to be evaluated tens of thousands of times. We used a few notable computational shortcuts. In particular, Equations 3 and 5 were combined to yield

$$g(t) = \sum_{j=1}^N k_j \int s(\tau) f_j(t - \tau) d\tau, \quad (8)$$

and the integrals were computed only once. Similarly, a set of functions $a(t)$, one for each simulated repetition of the stimulus, were computed only once for all iterations of the parameter search.

Error Measures

To evaluate the performance of the model, one wants to compare the set of real spike trains from repeated trials with a set of predicted spike trains from simulations with different noise waveforms. We started by parsing each set of spike trains into firing events as described by Berry and Meister (1998): the average firing rate was computed from a PSTH across trials (Figure 1B). This histogram was smoothed with a Gaussian filter whose width was adjusted to the time scale of variations in the firing rate, as determined from an auto-correlation of the spike trains (Berry et al., 1997). The resulting function generally showed sharp peaks separated by valleys of near zero firing rate. Formally, the boundaries between these firing events were established by finding minima v of the firing rate which were significantly lower than the neighboring maxima m_1 and m_2 such that $\sqrt{m_1 m_2}/v \geq 3$ with 95% confidence. The spikes between a pair of such boundaries were considered part of the same firing event. For each such firing event, we computed four response statistics of the distribution across trials:

$$T = \text{Average across trials of the time of the first spike,}$$

$$\begin{aligned}
 N &= \text{Average across trials of the spike number in the event,} \\
 V &= \text{Standard deviation across trials of the time of the} \\
 &\quad \text{first spike, and} \\
 S &= \text{Standard deviation across trials of the spike} \\
 &\quad \text{number in the event.}
 \end{aligned} \tag{9}$$

In this way a set of spike trains from multiple trials (Figure 1B) was converted into a single train of events (Figure 1C), each characterized by four numbers.

The next task was to evaluate the discrepancy between the real and simulated event trains. As illustrated in Figure 1C, some events in the predicted response clearly match those that occur in the actual response, but some events in each train may not have corresponding events in the other. To evaluate the quality of this correspondence, one wants to assess whether a given event in one train has a match in the other train, but also how well the four response variables (Equation 9) correspond between the actual and the predicted event. Suppose one had decided which events i in the real train are matched with which events j in the predicted train (by a method explained below), then we define the overall discrepancy between the two trains to be

$$E = e_T E_T + e_N E_N + e_V E_V + e_S E_S - e_M E_M. \tag{10}$$

This error measure contains four terms for the discrepancy in each of the event-specific properties

$$\begin{aligned}
 E_T &= \sum_{\text{matched event pairs } (i,j)} |T_i - T_j| \\
 E_N &= \sum_{\text{matched event pairs } (i,j)} |N_i - N_j| + \sum_{\text{unmatched events } i} N_i + \sum_{\text{unmatched events } j} N_j \\
 E_V &= \sum_{\text{matched event pairs } (i,j)} |V_i - V_j| \\
 E_S &= \sum_{\text{matched event pairs } (i,j)} |S_i - S_j|,
 \end{aligned} \tag{11}$$

where the symbols without primes indicate the properties of actual events, and the symbols with primes those of predicted events. The final term is proportional to the number of events matched between the two trains

$$E_M = \sum_{\text{matched event pairs } (i,j)} 1. \tag{12}$$

The constants e_x determine the relative importance assigned to different components of the error. For any given cell, we chose the particular values $e_T = 1/\bar{V}$, $e_N = 1/\bar{S}$, $e_V = 1/(2\bar{V})$, $e_S = 1/(2\bar{S})$, and $e_M = 2$, where a bar denotes the average value across all events of that neuron. The neuron's averaged timing jitter \bar{V} is a natural choice for scaling the errors in predicting the event time T ; in effect this compares the discrepancy between the predicted and actual response to the variation across trials of the actual response. In the same manner, the averaged spike number jitter \bar{S} serves to scale the errors in predicting the spike number N . The coefficients e_V and e_S are half as large as e_T and e_N because we wished to make fitting of the trial-to-trial variations V and S relatively less important than fitting the mean properties of an event T and N . Finally, the coefficient e_M is a "bonus" for matching two events, and its effects are discussed further below.

The above measure of error E (Equation 10) assumes that one knows which events i in the real train should be matched with which events j in the predicted train. We choose that set of matches which, in turn, produces the smallest error E . Fortunately, one does not need to inspect and evaluate all possible correspondences between the two trains, as long as one observes a natural restriction: that matches should not cross in time (Figure 1C). Two events in one train cannot be matched to events in the other train that occur in the opposite order. With this constraint, one can find the optimal correspondence between the event trains as follows (Victor and Purpura, 1996). Consider two event trains R and R' . At least one of the following three possibilities is true: the last event in R is un-

matched; the last event in R' is unmatched; or the last events in R and R' are matched to each other. Letting $E_{i,j}$ denote the error for the first i events of R and the first j events of R' , the three possibilities imply that

$$E_{i,j} = \min(E_{i-1,j} + e_N N_i, E_{i,j-1} + e_N N'_j, E_{i-1,j-1} + M_{i,j}), \tag{13}$$

where

$$\begin{aligned}
 M_{i,j} &= e_T |T_i - T'_j| + e_N |N_i - N'_j| + e_V |V_i - V'_j| + \\
 &\quad e_S |S_i - S'_j| - e_M
 \end{aligned} \tag{14}$$

is the error incurred in matching the last two events (i,j). The quantities $E_{i,j}$ can be viewed as a two-dimensional array. Starting with $E_{0,0} = 0$, this array can be filled recursively using Equation 13.

In practice, the array $E_{i,j}$ may have more than a thousand elements on each side, but not all of these need to be computed. For two events far apart in time—in particular if

$$|T_i - T'_j| > \frac{1}{e_T} (2e_N N_{\max} + e_M), \tag{15}$$

where N_{\max} is the largest number of spikes in an event in either event train—the error $M_{i,j}$ is so large that the possibility of a match need not be considered. Thus, one needs to evaluate only those $E_{i,j}$ for which the times of the last events meet Equation 15 and the array elements immediately adjacent to those. With this shortcut, the numerical effort grows proportionally to the total number of events, rather than its square, and the computation time for this matching step becomes negligible.

The Initial Guess and Optimization of the Model

A first guess for the various parameters can be obtained by correlating the spike train with the stimulus; this and other analytical properties of the spike generator will be derived in a subsequent paper (J.K., S. Smirnakis, and M.M., unpublished data). In brief, we computed the spike-triggered average stimulus waveform as well as the spike-triggered covariance about that average (de Ruyter van Steveninck and Bialek, 1988). The filter waveform was chosen as the first eigenvector of this covariance matrix. This analysis also yielded estimates of the threshold θ and the noise standard deviation σ_a . We further chose $B = \theta$, $\sigma_b = \sigma_a/\theta$, and $\tau_p = 0.2$ s (salamander and rabbit) or 0.02 s (cat).

Starting with this initial guess, the model was used to simulate a set of spike trains, with as many repeats as in the measured responses. Events were identified in both data sets and matched to each other as described above, yielding the error measure for this parameter set. Then we performed a search in the space of 20 parameters, repeating this simulation and evaluation at each step, to find a set that minimized the error. The search was implemented by Powell's method (Press et al., 1992), supplemented by simulated annealing (10,000 steps, reducing the "temperature" from 0.15 to 0.0005 in 10 geometric steps [Press et al., 1992]), which avoids getting trapped in local minima of the error surface. Minimization for a larger number of steps or by different methods did not significantly decrease the final error value.

The last coefficient of the error measure, e_M in Equation 10, was essential for the fitting process. It effectively encourages the matching of two nearby events in the two trains. Suppose there is a firing event with N_i spikes in the actual response, but the model predicts no event nearby, thus the contribution to the error is $e_N N_i$. Now we consider a small change in parameters, which leads to the prediction of an event close to the correct time, and ask whether the algorithm will match this to the actual event. Because the parameters have changed only slightly, the generator potential barely crosses threshold, and the predicted number of spikes is small, $N'_j \ll N_i$. If the two events were to remain unmatched, their contribution to the error would be $e_N(N_i + N'_j)$. If the two events were matched, the error from their spike numbers would be only slightly smaller, $e_N(N_i - N'_j)$. However, matching the two events introduces additional error terms in Equation 10 from comparisons of the event time, T , and the standard deviations S and V . These penalties will outweigh the benefits of matching the two events. In absence of a compensating reward, the new parameter set would be rejected, and the search would never explore promising regions of the parameter space.

Adding e_M eliminates this problem and allows a “budding” predicted event to be matched to an actual event.

Acknowledgments

This work was supported by NIH grants EY10020 to M.M. and EY10115, EY12196, and NS07009 to R.C.R. and P.R. We thank the members of our laboratories for comments on the manuscript. J.K. and M.M. designed and carried out experiments on salamander and rabbit retina, conceived and implemented the model, and wrote this article. P.R. and R.C.R. designed and carried out experiments on cat retina and LGN, and participated in their analysis and interpretation.

Received November 16, 2000; revised March 13, 2001.

References

- Bair, W., and Koch, C. (1996). Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput.* **8**, 1185–1202.
- Baylor, D.A., and Fettiplace, R. (1979). Synaptic drive and impulse generation in ganglion cells of turtle retina. *J. Physiol.* **288**, 107–127.
- Benardete, E.A., and Kaplan, E. (1997). 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. Natl. Acad. Sci. USA* **94**, 5411–5416.
- Buracas, G.T., Zador, A.M., DeWeese, M.R., and Albright, T.D. (1998). Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* **20**, 959–969.
- Clifford, C.W., and Ibbotson, M.R. (2000). Response variability and information transfer in directional neurons of the mammalian horizontal optokinetic system. *Vis. Neurosci.* **17**, 207–215.
- 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.
- de Ruyter van Steveninck, R.R., and Bialek, W. (1988). Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proc. R. Soc. Lond. B* **234**, 379–414.
- DeVries, S.H. (1999). Correlated firing in rabbit retinal ganglion cells. *J. Neurophysiol.* **81**, 908–920.
- DeVries, S.H., and Baylor, D.A. (1997). Mosaic arrangement of ganglion cell receptive fields in rabbit retina. *J. Neurophysiol.* **78**, 2048–2060.
- Enroth-Cugell, C., and Robson, J.G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *J. Physiol.* **187**, 517–552.
- Fohlmeister, J.F., Poppele, R.E., and Purple, R.L. (1977). Repetitive firing: a quantitative study of feedback in model encoders. *J. Gen. Physiol.* **69**, 815–848.
- French, A.S., and Korenberg, M.J. (1989). A nonlinear cascade model for action potential encoding in an insect sensory neuron. *Biophys. J.* **55**, 655–661.
- French, A.S., and Korenberg, M.J. (1991). Dissection of a nonlinear cascade model for sensory encoding. *Ann. Biomed. Eng.* **19**, 473–484.
- Hochstein, S., and Shapley, R.M. (1976). Linear and nonlinear spatial subunits in Y cat retinal ganglion cells. *J. Physiol.* **262**, 265–284.
- Humayun, M.S., de Juan, E., Jr., Weiland, J.D., Dagnelie, G., Katona, S., Greenberg, R., and Suzuki, S. (1999). Pattern electrical stimulation of the human retina. *Vision Res.* **39**, 2569–2576.
- Hunter, I.W., and Korenberg, M.J. (1986). The identification of nonlinear biological systems: Wiener and Hammerstein cascade models. *Biol. Cybern.* **55**, 135–144.
- Kara, P., Reinagel, P., and Reid, R.C. (2000). Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* **27**, 635–646.
- Kistler, W.M., Gerstner, W., and van Hemmen, J.L. (1997). Reduction of Hodgkin-Huxley equations to a single-variable threshold model. *Neural Comput.* **9**, 1015–1045.
- Knight, B.W. (1972). Dynamics of encoding in a population of neurons. *J. Gen. Physiol.* **59**, 734–766.
- Kondoh, Y., Arima, T., Okuma, J., and Hasegawa, Y. (1991). Filter characteristics of cercal afferents in the cockroach. *J. Comp. Physiol. A* **169**, 653–662.
- Korenberg, M.J., French, A.S., and Voo, S.K. (1988). White-noise analysis of nonlinear behavior in an insect sensory neuron: kernel and cascade approaches. *Biol. Cybern.* **58**, 313–320.
- 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.
- Marmarelis, P.Z., and Marmarelis, V.Z. (1978). *Analysis of Physiological Systems: The White-Noise Approach* (New York: Plenum Press).
- Meister, M., and Berry, M.J. (1999). The neural code of the retina. *Neuron* **22**, 435–450.
- Meister, M., Pine, J., and Baylor, D.A. (1994). Multi-neuronal signals from the retina: acquisition and analysis. *J. Neurosci. Methods* **51**, 95–106.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., and Flannery, B.P. (1992). *Numerical Recipes in C: The Art of Scientific Computing, Second Edition* (Cambridge: Cambridge University Press).
- 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.
- Reich, D.S., Victor, J.D., and Knight, B.W. (1998). The power ratio and the interval map: spiking models and extracellular recordings. *J. Neurosci.* **18**, 10090–10104.
- Reid, R.C., and Alonso, J.M. (1995). Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* **378**, 281–284.
- Reinagel, P., and Reid, R.C. (2000). Temporal coding of visual information in the thalamus. *J. Neurosci.* **20**, 5392–5400.
- 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.
- 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. (1995). Response dynamics and receptive-field organization of catfish ganglion cells. *J. Gen. Physiol.* **105**, 795–814.
- Sakai, H.M., Naka, K., and Korenberg, M.J. (1988). White-noise analysis in visual neuroscience. *Vis. Neurosci.* **1**, 287–296.
- Sakuranaga, M., and Naka, K. (1985). Signal transmission in the catfish retina. III. Transmission to type-C cell. *J. Neurophysiol.* **53**, 411–428.
- Sakuranaga, M., Ando, Y., and Naka, K. (1987). Dynamics of the ganglion cell response in the catfish and frog retinas. *J. Gen. Physiol.* **90**, 229–259.
- Shadlen, M.N., and Newsome, W.T. (1994). Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* **4**, 569–579.
- 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.
- Usrey, W.M., Reppas, J.B., and Reid, R.C. (1999). Specificity and strength of retinogeniculate connections. *J. Neurophysiol.* **82**, 3527–3540.
- Victor, J.D. (1987). The dynamics of the cat retinal X cell centre. *J. Physiol.* **386**, 219–246.
- 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.

Victor, J., and Shapley, R. (1980). A method of nonlinear analysis in the frequency domain. *Biophys. J.* *29*, 459–483.

Victor, J.D., and Purpura, K.P. (1996). Nature and precision of temporal coding in visual cortex—a metric-space analysis. *J. Neurophysiol.* *76*, 1310–1326.

Vinje, W.E., and Gallant, J.L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* *287*, 1273–1276.

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.