

Multineuronal Firing Patterns in the Signal from Eye to Brain

Mark J. Schnitzer¹ and Markus Meister^{2,*}

¹Biological Computation and Theoretical Physics

Research Department

Bell Laboratories

Lucent Technologies

Murray Hill, New Jersey 07974

²Molecular and Cellular Biology Department

Harvard University

Cambridge, Massachusetts 02138

Summary

Population codes in the brain have generally been characterized by recording responses from one neuron at a time. This approach will miss codes that rely on concerted patterns of action potentials from many cells. Here we analyze visual signaling in populations of ganglion cells recorded from the isolated salamander retina. These neurons tend to fire synchronously far more frequently than expected by chance. We present an efficient algorithm to identify what groups of cells cooperate in this way. Such groups can include up to seven or more neurons and may account for more than 50% of all the spikes recorded from the retina. These firing patterns represent specific messages about the visual stimulus that differ significantly from what one would derive by single-cell analysis.

Introduction

If a neuronal population contains N spiking cells, how diverse is the set of possible messages these neurons can convey by their firing? The answer depends on the coding scheme used to represent the messages. For example, in the classical “labeled line” conception of sensory coding, the meaning of a spike is determined entirely by the identity of the neuron firing the spike, regardless of the activity in other neurons. In that case, each neuronal spike train represents an independent channel of communication, and the number of such channels conveyed by the population grows proportionally with N . Alternatively, one might imagine that the firing of one neuron changes the meaning of spikes from another neuron. In the extreme case, each pattern of joint activity among N neurons has a different meaning; then the total number of distinct messages grows exponentially with N . Thus, a distributed code that employs the concerted activity of multiple neurons would enjoy much greater representational power than a scheme in which individual neurons transmit messages independently. This comes at a cost of more action potentials, since any given message may require firing by multiple neurons rather than by a single cell (Meister, 1996). Do real neurons encode information with such a concerted scheme? This fundamental question has been difficult

to address, because three important components must all come together.

First, one must record from many neurons simultaneously to even get access to their firing patterns. Until recently this was a formidable obstacle, but technical advances over the last decade have made multielectrode recording increasingly widespread (Buzsaki et al., 1989; Wilson and McNaughton, 1993; Kovacs et al., 1994; Meister et al., 1994; Nordhausen et al., 1996; Della Santina et al., 1997; Chapin and Nicolelis, 1999).

Second, one must recognize those firing patterns that are plausible candidates for concerted coding. A brute force inspection of every possible combination of spikes among N neurons is prohibitive for even moderate N (Gerstein and Aertsen, 1985). Many studies have so far been limited to pairwise comparisons among neurons, for example through analysis of crosscorrelation functions. Identifying higher-order correlations among multiple neurons is a statistical challenge (Gerstein et al., 1978; Gerstein and Aertsen, 1985; Abeles and Gerstein, 1988; Frostig et al., 1990). We have now developed an efficient algorithm that can detect concerted activity of arbitrarily large groups of neurons while circumventing the combinatorial explosion of the brute force approach.

Finally, one must be able to interpret what messages are conveyed by neural firing, to test whether concerted firing patterns play a special role. The opportunities for this are greatest at the sensory or motor periphery of the nervous system, simply because the variables of interest to these neurons—sensory stimuli or muscle movements—are directly observable. Here we analyze how ganglion cells in the retina encode visual stimuli. The output of 50 or more of these neurons can be recorded simultaneously with an electrode array, while visual stimuli are projected onto the photoreceptor layer (Meister et al., 1994). The goal is to determine whether concerted patterns of action potential firing by these neurons contribute to the representation of visual input.

Prior studies in several species have shown that nearby retinal ganglion cells have a strong tendency to fire in approximate synchrony (Arnett, 1978; Arnett and Spraker, 1981; Johnsen and Levine, 1983; Mastronarde, 1989; Meister et al., 1995; Meister, 1996; Brivanlou et al., 1998; DeVries, 1999; Usrey and Reid, 1999). Pairwise correlation functions between two spike trains showed that this synchrony occurs on three different time scales (Mastronarde, 1989; Brivanlou et al., 1998). Pharmacological analysis in the salamander retina suggested synaptic circuits that might produce these different modes (Brivanlou et al., 1998). Electrical gap junctions likely couple pairs of ganglion cells that fire with tight synchrony (delays <1 ms). Shared excitation from an amacrine cell, again via electrical junctions, probably induces synchrony on an intermediate time scale (10–25 ms). Shared photoreceptor input conveyed through divergent interneuron pathways via chemical synapses leads to broad synchrony (40–100 ms) between two ganglion cells.

Although such correlations alone are not evidence for a concerted code, they merit further study because they

*Correspondence: meister@biosun.harvard.edu

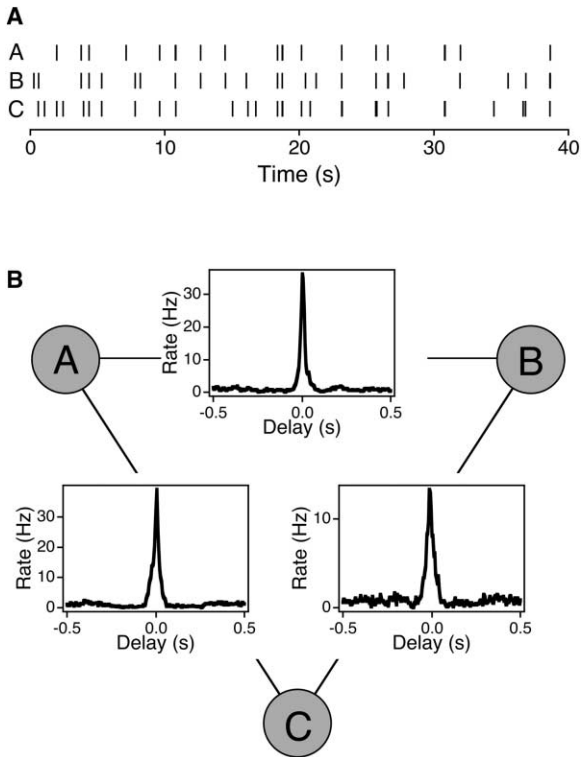


Figure 1. An Example of Three-Way Correlations among Retinal Ganglion Cell Spike Trains

(A) Brief segment of spike trains from three salamander retinal ganglion cells recorded in darkness.

(B) Pairwise crosscorrelograms with a sharp central peak show that each pair of cells has a tendency to fire synchronously. Yet, this pairwise analysis misses the frequent triplets of spikes that occur synchronously in all cells, as seen on direct inspection of (A).

might form the basis for such a scheme. We begin by testing whether these correlations are limited to pairwise interactions among ganglion cells or extend over larger groups. To do this, we introduce an information-theoretic algorithm that can identify arbitrarily large groups of cells engaged in concerted firing. This allows us to quantify the prevalence of concerted activity and to determine which firing patterns are candidates for a concerted code. We then inspect what messages these firing patterns convey about the visual stimulus, and show that these differ greatly from what one would expect if ganglion cells acted independently. Thus, groups of ganglion cells appear to engage in a concerted code for vision, and their output cannot be fully deciphered without distinguishing between synchronous and non-synchronous activity.

Results

An Algorithm to Find Groups of Synchronized Neurons

As reported previously (Mastronarde, 1989; Brivanlou et al., 1998; Usrey and Reid, 1999), pairs of nearby retinal ganglion cells had a strong tendency to fire in synchrony, even when recorded in complete darkness (Figure 1). Direct inspection of the spike trains in such experiments

often showed cases where many cells fired all at the same time (Figure 1A). This suggests that the pairwise correlations (Figure 1B) are merely the trace of synchronization among larger groups of neurons. To better understand synchronized activity, one would like to find such groups of cells that fire together more often than expected by chance, and that also do so frequently enough to contribute substantially to visual signaling.

The search for such stereotyped patterns has a close connection to data compression (Storer, 1988). Suppose, for example, that a particular group of cells always fired together. Then one could simply replace all their spike trains in the data set by a single spike train, without loss of information. This would yield a more compact representation of the neural recordings. The algorithm we used to find firing patterns operates on this same data compression principle (Figure 2A and Experimental Procedures). It searches for a pair of cells that often fire together and recodes those events in a new symbolic spike train, using a single spike for each pair of the original spikes. Then, the same procedure repeats on the new data set, including the symbolic spike train. After each round of iteration, fewer spikes are needed to represent the data, and from the amount of compression one can evaluate the significance of the most recently identified firing pattern (Figure 2B). This requires not only that the pattern occur frequently, but also that it occur more frequently than expected (Experimental Procedures). Iteration stops when no further compression is possible. One then simply evaluates all symbolic spike trains created by the algorithm (Figure 2B) to see what groups of neurons they represent.

The procedure's principal strength is that it can find arbitrarily large groups of coordinated neurons—by combining various symbolic spike trains—while requiring only a pairwise comparison of objects in each round of iteration. This avoids the dreaded combinatorial explosion, because the computational effort scales only quadratically with the number of neurons. A corollary is that the algorithm will recognize a large group of synchronized neurons only if some smaller subgroup was also deemed to be significant. In this sense, the method is conservative. Another important aspect is that the assignment of cells to patterns is not exclusive: one neuron can participate in several different groups (Figure 2B). This emerged as a useful feature, because such behavior occurs frequently in the retinal output signals.

Synchronized Firing in the Retina

We applied this method to search for synchrony in spike trains of retinal ganglion cells from tiger salamanders and rabbits. Synchrony between two spikes was defined as a time delay < 25 ms. This choice was made specifically to capture the “intermediate width” correlations among ganglion cells, which are the most prominent form of synchrony and are thought to arise in the inner retina (Mastronarde, 1989; Brivanlou et al., 1998). Initially, we inspected spontaneous activity in darkness. Our algorithm identified many groups of cells that fired synchronously. To quantify the prevalence of this concerted firing, we determined the fraction of all recorded spikes that occurred in multineuron firing patterns. In

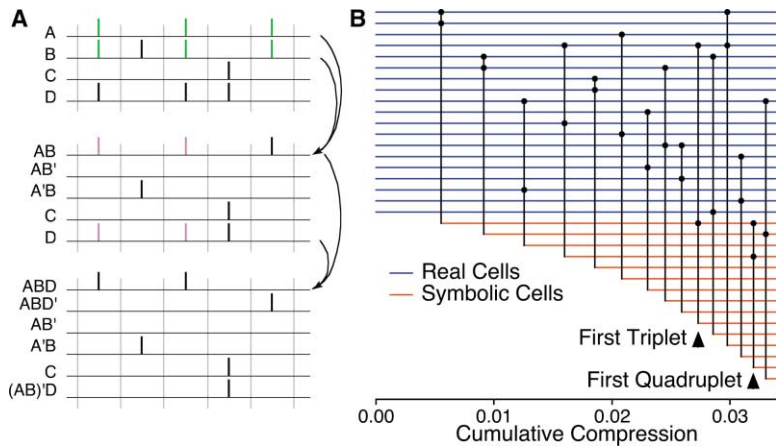


Figure 2. An Algorithm to Detect Multineuronal Firing Patterns

(A) Summary of the procedure. Top: schematic spike trains from four neurons (A, B, C, and D). Time runs left to right, divided in discrete bins. Middle: in the first round, the algorithm recognizes that cells A and B often fire in the same time bin and encodes each such pair of spikes with a single spike from a newly created symbolic cell AB. The events where either A or B fires alone are retained by the symbols AB' and A'B. Bottom: in the second round, the algorithm recognizes that symbols AB and D often occur synchronously and encodes these events with a single spike from symbolic cell ABD. This procedure continues until no more combinations can be found that satisfy the information theoretic criterion of significance.

(B) A graph illustrating progression of the iterative search procedure on a subset of retinal ganglion cell spike trains (salamander, darkness). Each horizontal line corresponds to a real cell or symbolic cell. Each vertical line indicates the combination of the two cells marked by a dot into a new symbol. Its horizontal offset indicates how much this definition contributed to compressing the data set, as a fraction of its original size. Arrows mark the identification of the first synchronized group containing three cells (the triplet of Figure 1) or four cells.

various retinæ this fraction was 16% (rabbit, 16 recorded cells), 17% (salamander, 23 cells), 43% (salamander, 44 cells), and 60% (salamander, 39 cells), and it increased the more ganglion cells were recorded by the electrode array. These numbers likely underestimate the actual contributions of activity in multineuronal firing patterns. First, the electrode array records from only a fraction of the ganglion cells in its vicinity ($\sim 15\%$; Meister et al., 1994), leading to a systematic underestimate of the size of groups, as some of the partners in the group may not be recorded (Experimental Procedures). Further, the search algorithm is not exhaustive and may have missed some spiking groups even among the cells whose activity we did record. Thus, the results presented here are conservative estimates.

To estimate the significance for the identified firing patterns, we measured how frequently a pattern appeared compared to the frequency expected if the participating cells were all firing independently of each other. The “correlation index” (Meister et al., 1995) is the ratio between the observed frequency of synchronous firing and that expected by chance (Experimental Procedures). Its distribution is shown in Figure 3 for synchronous groups of size 2 and 3. The typical synchronous spike pair occurred ~ 10 times more frequently than

expected, whereas the typical synchronous triplet occurred ~ 100 times more frequently than expected. For some triplets this ratio was 10^3 or more, and generally the correlation index was somewhat stronger under visual stimulation. As suggested by the example in Figure 1, the strength of correlation among groups of three cells was not simply explained by the underlying pairwise correlations: Triplets of spikes typically occurred 10 times more frequently than expected, assuming that the most significant pair and the single cell composing the triplet fired independently (Figure 3, inset).

Altogether, synchronous firing seems to account for half the retinal output or more, and it is clearly distinct from chance coincidences of spikes. Thus, it is worth studying these patterns further.

Properties of Synchronous Spiking Groups

As an example, we will examine in depth the responses in a representative population of 44 salamander ganglion cells, whose activity was recorded first in darkness and then under visual stimulation with a randomly flickering checkerboard (Meister et al., 1994, 1995). Many synchronized groups were identified in this population: 143 groups during spontaneous firing in darkness, and 99 groups under visual stimulation. Groups of two cells

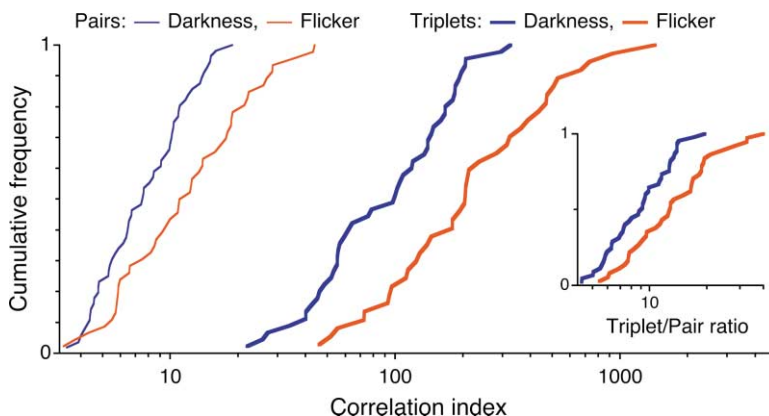


Figure 3. Strength of the Correlation in Spike Pairs and Triplets from Salamander Retina

A cumulative histogram of the correlation index, the ratio of the number of concerted action potentials from a group to the number expected by chance if the member cells fired independently of each other (Equation 10); note the logarithmic axis. Inset: the ratio of the number of triplets ABC to the number expected by chance if the pair AB and the cell C fired independently of each other (see Experimental Procedures).

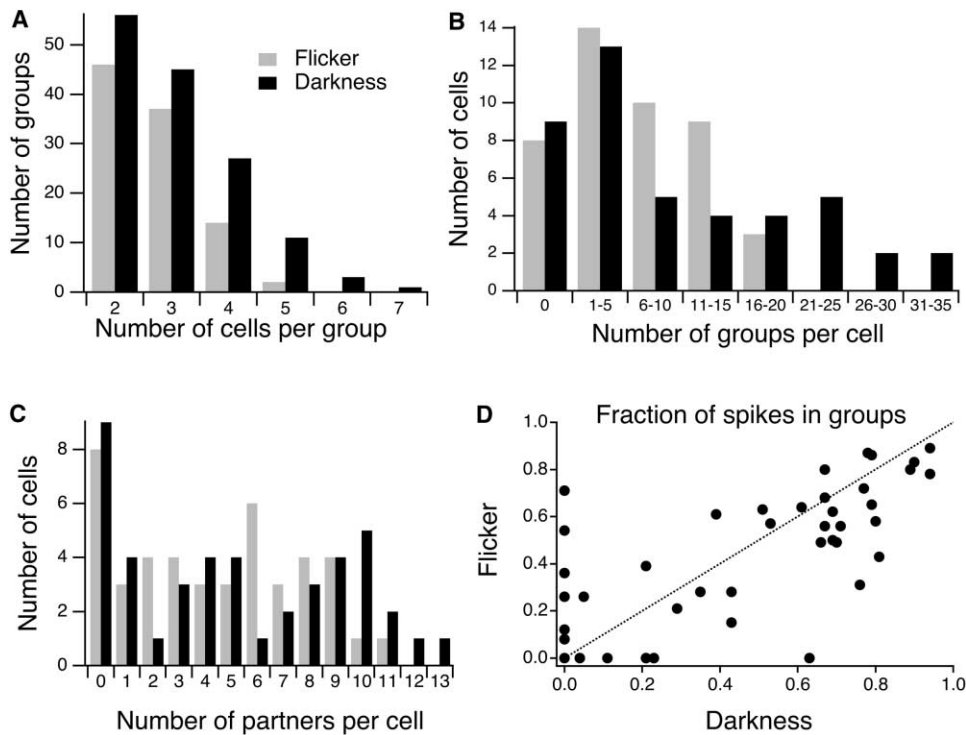


Figure 4. Summary Statistics of Synchronous Groups in a Ganglion Cell Population from Salamander Retina
 (A) Histogram across groups of the number of cells in the group, observed in darkness (solid bars) or under stimulation with a flickering checkerboard (shaded bars).
 (B) Histogram across cells of the number of groups in which a cell participates.
 (C) Histogram across cells of the number of partners with which a cell participates in groups.
 (D) The fraction of spikes fired in groups for each cell, with the value under flicker stimulation plotted against the value in darkness.

accounted for less than half the firing patterns, and larger groups were observed with decreasing frequency, up to groups of seven (Figure 4A). This confirms the need for a search method that is not restricted to pairs.

The number of different groups engaged in synchronous firing greatly exceeded the number of cells recorded in the population, so individual neurons clearly participated in more than one firing pattern. Especially in darkness, some cells were active in many groups, sometimes over 20 (Figure 4B). Other neurons participated in zero or only a few groups. The median number of groups per cell was 5.5. This raises the possibility that an individual ganglion cell may help transmit a variety of distinct visual messages, depending on which other cells spike simultaneously. In principle, a small number of neurons would be sufficient to construct many different groups: for example, four neurons could form nine different synchronous groups. Thus, it is interesting to examine the actual number of distinct partner cells that a neuron engages within all its groups. This was found to vary over a broad range, up to >10 partners (Figure 4C). Still, a distinct subset of cells (~20%) did not fire in groups at all.

Salamander ganglion cells can be classified into four major types based on their visual response properties (Warland et al., 1997), and these types contributed very differently to synchronous firing: within the “fast OFF” and “slow OFF” types, the proportion of neurons firing in groups was 80%, among “weak OFF” cells 30%, and

among ON cells only 10%. Apparently, certain types of retinal ganglion cell are less influenced by the lateral networks that underlie synchronization (Brivanlou et al., 1998; DeVries, 1999). For these neurons, visual signaling may perhaps be understood based on single-cell responses alone.

The randomly flickering checkerboard is a strong stimulus for retinal ganglion cells, and one might expect this to dramatically alter the patterns in which neurons fire. In general, the number of significant groups decreased under visual stimulation (Figure 4A), but the remaining ones were similar to those in darkness. In the above sample population, 37% of the groups encountered under flicker stimulation also fired together in darkness. For individual neurons, the fraction of spikes in groups was highly correlated in darkness and flicker, although a few cells participated in no groups under one condition or the other (Figure 4D). One concludes that synchronous firing in groups is prominent with and without visual input, and that some of the networks that synchronize neurons in darkness are also active under visual stimulation.

What is the spatial arrangement of retinal ganglion cells engaged in synchronous firing? Prior analysis of pairwise correlations has shown that the correlation index decreases with distance between two cells, by a factor of 1/e in ~200 μm (Meister et al., 1995). Accordingly, we found that cells participating in a group tend to be close to each other, though not necessarily nearest

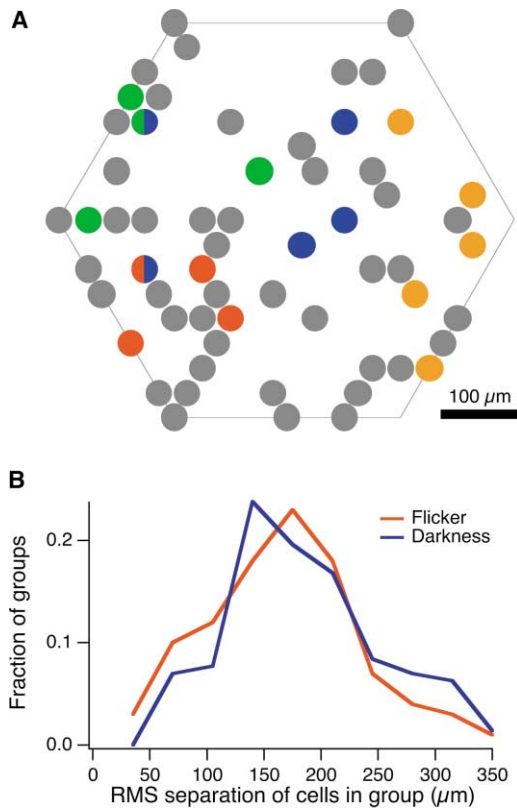


Figure 5. The Spatial Relationships of Synchronized Groups of Neurons

(A) The locations on the retina of 44 ganglion cells recorded simultaneously, with an outline of the hexagonal multielectrode array. Four of the ganglion cell groups found in darkness are identified by a common color. Cells filled with two different colors participated in both groups.

(B) The separation of cells in a group was measured as their root mean square distance and histogrammed over all groups recorded in darkness or under flicker stimulation.

neighbors (Figure 5A). Often, a group of cells was interspersed with other neurons that did not participate in the same pattern of firing. We quantified the spatial extent of a group by the root mean square distance between its component cells. This ranged from $\sim 50 \mu\text{m}$ to $\sim 350 \mu\text{m}$ for different groups (Figure 5B). The spatial extent was similar for groups found in darkness ($174 \pm 75 \mu\text{m}$, mean \pm SD) and under visual stimulation ($150 \pm 65 \mu\text{m}$). The lateral networks underlying these correlations apparently extend for a few hundred micrometers, but connect to ganglion cells within that range selectively rather than indiscriminately.

Visual Messages Transmitted by Synchronous Spiking Groups

We now consider what information a synchronous firing pattern conveys to the brain about the visual stimulus, and compare this to the visual meaning of such a firing pattern if it arose by chance among ganglion cells responding to the stimulus independently.

To provide the retina with a rich visual ensemble, we stimulated it with a randomly flickering checkerboard pattern (Experimental Procedures; Meister et al., 1994).

In this “white noise” display, the light intensity is modulated randomly and independently in time, space, and the color dimension. From a ganglion cell’s spike train during a long stretch of such flicker stimulation, one can compute the spike-triggered average stimulus (STA). This corresponds to the average visual sequence—equivalent to a short movie clip—that precedes a spike from the neuron (Experimental Procedures; Meister et al., 1994). To first order, this is the visual message that the neuron conveys to the brain: if an observer at the other end of the optic nerve records a spike from this cell, in the absence of additional information, the best estimate of the preceding visual stimulus is equal to the neuron’s STA. This analysis easily extends to groups of synchronized neurons: to find the message conveyed by a firing pattern, one calculates the average stimulus triggered on the occurrence of synchronous spikes from all cells in the group (Ghose et al., 1994; Meister et al., 1995; Dan et al., 1998).

Figure 6 illustrates the STA for three retinal ganglion cells, and also for their synchronous firing pattern. Note that the STA is a function of space, time, and color (Equation 16), which can be approximated usefully as a product of a spatial profile, a time course, and a color spectrum (Equation 17). The three ganglion cells in Figure 6A are very similar in the time course and in the spectral sensitivity of their STA, reflecting common characteristics of the “strong OFF” functional type of ganglion cells (Meister et al., 1995). Specifically, the time course shows a characteristic dip in the intensity at negative times, so a spike from such a neuron reports that a dimming occurred $\sim 100 \text{ ms}$ earlier. The spectral sensitivity is largest for the red channel of the monitor, because the salamander retina is dominated by a red-sensitive cone. The three cells differ in the spatial profile of their receptive field, which indicates that their dendrites cover different but overlapping regions of the retina. The STA for the synchronous firing events involving all cells resembles the individual STAs in the temporal and spectral dimensions. However, the receptive field profile of this firing pattern differs from that of the individual cells: it occupies a small region approximately at the intersection of their individual receptive fields.

The neurons participating in a synchronous group tended to be very similar in their temporal integration and spectral sensitivity, but differed in their spatial receptive fields (e.g., Figure 6A). Among cells in a group, the relative variation (Equation 21) was $19\% \pm 12\%$ (mean \pm SD over groups) for the time course of the STA ($a(t)$ in Equation 17), only $6\% \pm 7\%$ for the spectral sensitivity (c_k), but $73\% \pm 27\%$ for the receptive field (bx). This is consistent with the observation that synchronized firing occurs primarily among ganglion cells of similar functional types, rather than across a mixture of types (Meister et al., 1995; DeVries, 1999). Thus, if synchronous spikes play any special role in communicating the visual stimulus, it is likely in the spatial domain, by specifying where a stimulus occurred, not when it occurred or what color it had. Accordingly, we focused further analysis on the spatial receptive fields of firing patterns.

In this context, it is useful to have a benchmark expectation for the receptive field of a firing pattern, as measured by the spatial profile of its STA. What should one

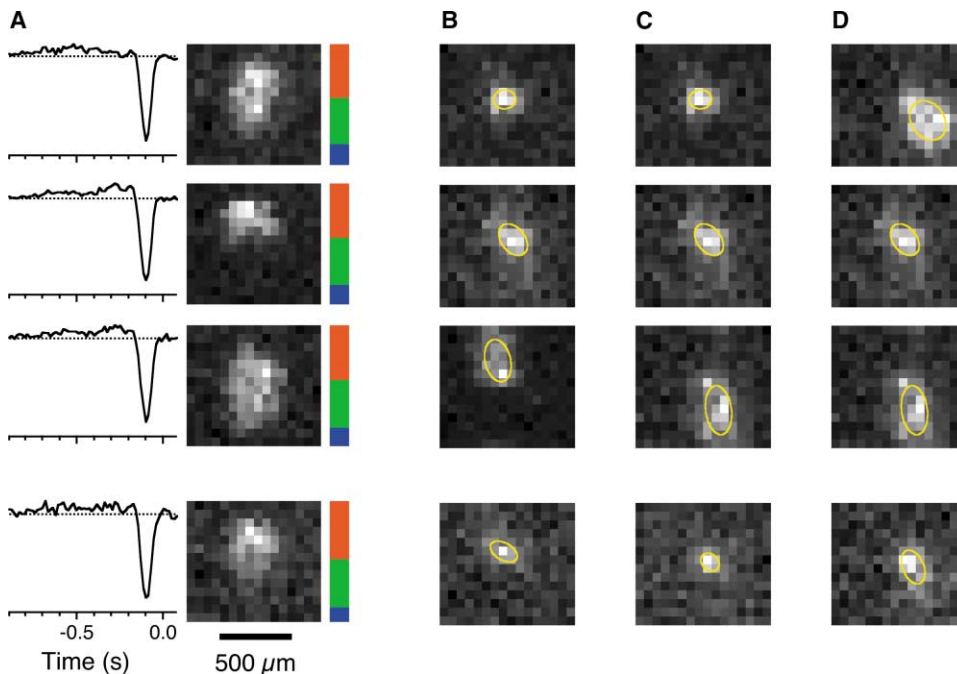


Figure 6. The Spike-Triggered Average Stimulus of Cells and Synchronous Groups

(A) The STA for three ganglion cells (top three panels) and for the synchronous spike triplet among those three cells (bottom panel). Each panel illustrates the time course of the STA (left, $a(t)$ in Equation 17), the spatial profile (middle, $b(x)$), and the relative sensitivity to the three color channels (right, c_i).

(B) The receptive field profiles, $b(x)$, of three individual cells and of their synchronous spike pattern (bottom), displayed as in the middle column of (A). The ellipse is the contour of the Gaussian shape used to fit the receptive field profile, at 1 standard deviation from the center.

(C and D) Receptive fields for two more triplets of cells and their synchronous spikes, displayed as in (B). Note that some cells participate in several triplets.

expect for this receptive field if the individual ganglion cells are not engaged in any concerted function, but simply modulate their firing rate independently of each other in response to the stimulus? Under this null hypothesis, synchronous spikes occur only because the neurons are individually modulated to fire faster at around the same time. One can make a simple prediction for the STA of synchronous spikes in the special case where the neurons have nonoverlapping receptive fields. For example, two otherwise identical OFF cells with disjoint receptive fields will fire synchronously if the random checkerboard has dark values simultaneously in those two distinct regions. After averaging the stimulus over many such events, one will obtain a receptive field profile with two lobes that looks like the sum of the two individual receptive fields. This relationship can be generalized to cells with overlapping receptive fields (Equation 32), showing that, under a broad range of conditions, the STA for synchronous spikes among independent neurons is proportional to the sum of their individual STAs.

Under the null hypothesis of independent visual signaling, one would therefore expect the receptive field of a firing pattern to approximate the union of the individual cells' receptive fields, which is of course always larger than the individual receptive fields. Counter to this prediction, we found that the receptive field of a firing pattern was generally smaller than that of the individual participating neurons, illustrated in Figures 6B–6D with

examples of spike triplets. To analyze this further, each STA profile was fitted with a two-dimensional Gaussian (Figures 6B–6D), and we measured the receptive field radius by the standard deviation of that Gaussian (Equation 20). Figure 7A compares the receptive field radius of each firing pattern to that of the average participating neuron. We found that the pattern's receptive field was nearly always smaller than that of the average component cell, and certainly smaller than the union of their receptive fields. This was true irrespective of the size of the group. For $\sim 60\%$ of the groups inspected, the receptive field radius was even smaller than that of the smallest receptive field among the component cells.

We inspected receptive field sizes by a second method that does not rely on Gaussian fitting. In this approach, we measured the area of a receptive field in which the magnitude of the STA exceeded a given threshold. This was done for the synchronous firing pattern and for the individual participating cells. We compared the spatial union of the regions above threshold for the individual cells with the area above threshold for the firing pattern; this was repeated for several different threshold values. The method confirmed that the receptive field of a synchronous group was considerably smaller in area than the sum of the receptive fields of its component cells (Figure 7B). We conclude that the visual message conveyed by a synchronous spiking group is different from what one would expect given the visual responses of the component cells taken individu-

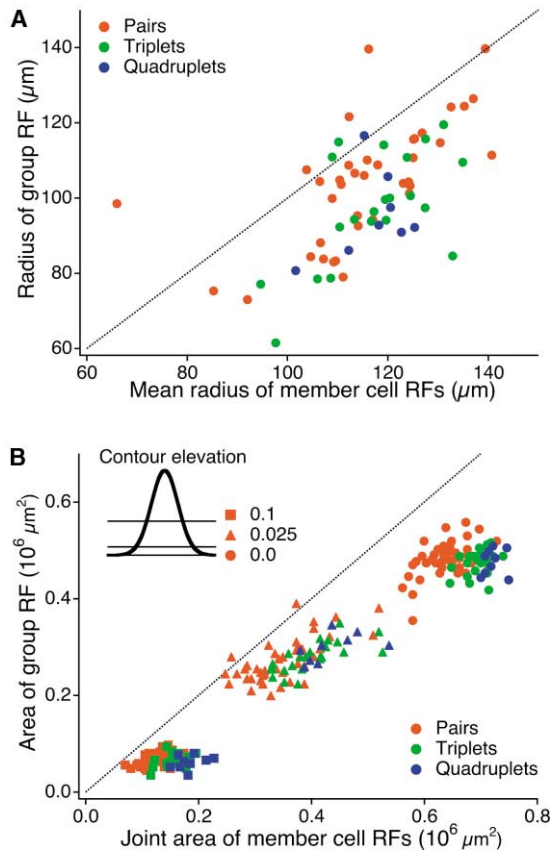


Figure 7. The Receptive Field Sizes of Synchronous Groups

(A) The radius of the visual receptive field of a synchronous group of ganglion cells is compared to the average radius of the individual receptive fields of the members of the group. Results are shown separately for ganglion cell pairs, triplets, and quadruplets. The line represents equality. Radii are determined from the Gaussian fits (Figures 6B–6D and Equation 20).

(B) The area covered by the receptive field of a synchronous group of ganglion cells is compared to the area of the union of their individual receptive fields. In each case, the area is that enclosed by a contour line of the receptive field profile, $b(x)$. To test the generality of the relationship, the contours were measured at a series of different elevations of the receptive field profile, indicated by different symbols.

ally. Specifically, the receptive fields of firing patterns are more localized in space and can convey spatial detail at a resolution finer than expected from receptive fields of single ganglion cells.

Discussion

The Origin of Synchronous Firing Patterns

The observation of synchronous firing among retinal ganglion cells in the absence of visual stimulation (Figure 1) suggests that they share a source of presynaptic input which is spontaneously active even in darkness (Mastrorarde, 1983; Meister et al., 1995). The present study focused on synchrony on the short and intermediate time scales (defined here by <25 ms spike interval), thought to be caused by shared input to ganglion cells via electrical synapses (Brivanlou et al., 1998). This

places the shared presynaptic neuron in the inner retina: an amacrine cell or possibly another ganglion cell (Figure 8A). Because those correlations are so short-lived, this cell likely produces action potentials, or at least fast transients, even during spontaneous activity. If this interneuron branches out to several ganglion cells, then those will all be excited together and become recognized as a synchronously firing group in our analysis.

Under visual stimulation, we found many of the same ganglion cell groups as in darkness, suggesting that the interneurons that synchronize ganglion cells in darkness also do so under visual stimulation (Figure 8B). In this view, the receptive field of a synchronized spike pattern among ganglion cells is simply the receptive field of the underlying interneuron. An individual ganglion cell often participates in more than one pattern of firing (Figure 4B), suggesting that it is driven by several such interneurons. Then the ganglion cell's receptive field profile, as measured by the STA, is the average of the receptive fields of its presynaptic interneurons (Figure 8B and Experimental Procedures). This can explain why the receptive fields of firing patterns generally had a smaller extent than those of individual ganglion cells (Figures 6 and 7). Note also that the receptive field size of a firing pattern does not depend on the number of ganglion cells in the group (Figures 7A and 7B). For example, in the thresholding analysis of Figure 7B, the receptive field area for triplets was $95\% \pm 20\%$ (mean \pm SD) that of the underlying pairs. Pairs, triplets, quadruplets, and higher-order patterns all have a similar receptive field size, which is considerably smaller than that of a single cell. This is consistent with the notion that any firing pattern among ganglion cells reflects the firing of a presynaptic interneuron, and those neurons all have similar receptive fields.

If one were to analyze the ganglion cell responses one neuron at a time, one could certainly derive an understanding of how each neuron responds individually to the visual stimulus. One even expects that nearby neurons should occasionally fire together because of overlap in their spatial receptive fields. However, the prediction of these synchronized firings would be grossly flawed. For example, a popular framework for visual responses from individual ganglion cells is the "linear-nonlinear" (LN) model (Hunter and Korenberg, 1986; Meister and Berry, 1999; Chichilnisky, 2001): it assumes that the light intensity gets pooled over space and integrated over time with a linear weighting function (L), and the result gets converted to the neuron's firing rate by a generally nonlinear function (N). Such a model can accommodate many key features of visual neurons: an arbitrary receptive field profile; an arbitrary temporal integration function and spectral sensitivity at each point in the receptive field; and nonlinearities such as rectification and saturation in the firing response. Using any model in this class for the single-cell responses, one can show that the receptive field profile of a joint firing pattern is the average of the component ganglion cell receptive fields, and thus should be larger in extent than that of the individual neurons (Figure 8C and Equation 33). This is very different from what we observed (Figures 6 and 7).

Synchronous firing patterns account for up to 60% of all the spikes recorded by the electrode array and up

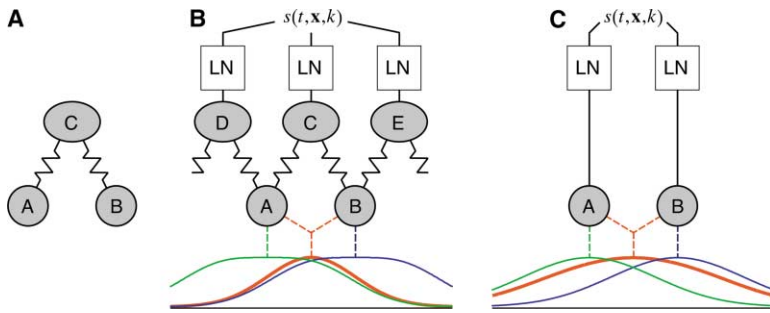


Figure 8. Circuitry Contributing to Synchronous Firing among Ganglion Cells

(A) Synchronous firing of two ganglion cells (A and B) in darkness may be caused by coupling to a spontaneously active amacrine cell or other interneuron (C).

(B) Under visual stimulation, the interneurons (C, D, and E) respond to the stimulus, here via an LN mechanism (Equation 22). Each interneuron is taken to have a Gaussian receptive field profile. For each of the ganglion cells (A and B), the receptive field profile (thin bottom curves) is the average of that of its input neurons (Equation 36). The joint spikes

between two ganglion cells have the narrow receptive field (bold bottom curve) of their shared interneuron.

(C) Here the ganglion cells (A and B) respond to the stimulus independently, each by an LN mechanism (Equation 29) with a Gaussian receptive field (thin bottom curves). Coincidences between the ganglion cells occur when the stimulus excites each one. These events have a broad receptive field (bold bottom curve) corresponding to the average of the individual ganglion cell receptive fields (Equation 33).

to 90% for certain individual cells. Moreover, their response properties are not what one might predict based on single-unit analysis, so it is important to recognize the firing patterns when interpreting the activity of this neural population.

The Interpretation of Synchronous Firing Patterns

To extract visual information from these responses optimally, one should assign a separate message to each firing pattern, rather than to each individual neuron. For example, in Figures 6B–6D the same cell participates in all three spike triplets, but its spikes have a different receptive field in each of those combinations and thus convey a different message about the location of the stimulus. Because the receptive fields of firing patterns are small and there are many more such patterns than there are individual cells, this way of decoding the optic nerve signals could reveal greater spatial detail than one treating individual neurons as the fundamental elements. Such an enhancement of spatial information has indeed been achieved in decoding the synchronous firing of LGN neurons in the cat (Dan et al., 1998).

In the wiring diagram of Figure 8B, the visual scene is encoded at the spatial resolution of the presynaptic interneurons, without requiring a dedicated “labeled line” for each of these neurons in the optic nerve. The signal of each interneuron is instead “labeled” by a synchrony tag between several optic nerve fibers. The precise *relative* timing of spikes from different ganglion cells identifies them as a group carrying its own message about a distinct visual feature. The *absolute* timing or frequency of that spike pattern encodes when or how much of that feature occurred in the visual stimulus. The exact nature of that relationship between stimuli and spike groups remains to be spelled out, but one can think about neural coding by spike groups using the same kinds of models previously applied to single spike trains (Meister and Berry, 1999).

An important condition for such concerted coding is that the average firing rate of ganglion cells must be low enough to avoid spurious coincidences that could be interpreted as synchrony tags. This is indeed the case, in both amphibian and mammalian retina (Meister, 1996). The retina can afford these low average firing rates because the temporal variation in the stimulus,

after filtering through the outer retina, is rather slow, on the time scale of ~ 100 ms (Warland et al., 1997). In effect, the excess temporal bandwidth that fast-spiking ganglion cells offer over slow photoreceptors can be converted to additional spatial bandwidth by the use of a synchrony code.

Is such a concerted coding scheme actually used to productive ends by the visual system? Studies of retinal responses alone cannot provide the answer, but there are additional pieces of suggestive evidence (Meister and Berry, 1999). For one, virtually all vertebrate eyes have considerably more photoreceptors than optic nerve fibers, and thus could support a finer spatial resolution than one would estimate from the array of ganglion cells. Second, the optic nerve is the narrowest part of the visual system; nowhere else in the brain is the visual scene represented by as small a population of neurons. Constraints on the thickness of this cable arise in part from the need for eye movement, which pulls and bends the nerve. Under these conditions, a coding scheme that conveys more distinct messages than there are fibers would indeed be very attractive.

The neural hardware in the early visual system seems well suited to supporting a synchrony code (Usrey and Reid, 1999). Synapses in the mammalian retino-geniculate pathway are exceptionally strong, unlike intracortical synapses. In fact, a single retinal ganglion cell can trigger an LGN relay neuron in almost 1:1 fashion (Cleland et al., 1971), and thus conveys its spike train reliably to the visual cortex. Groups of synchronous spikes will be preserved in this process, because there is little temporal dispersion. Neurons participating in a synchronous group tend to be of the same cell type and located near each other on the retina; thus, they will have axons of similar length and diameter, and therefore similar conduction times to the target. In fact, direct measurements of conduction times in the cat suggest that the total latency from a retinal ganglion cell through the LGN to the cortex varies by <3 ms among cells of the same functional class (Cleland et al., 1976). Once the signals reach the visual cortex, the neural bottleneck is passed, and a much larger neuronal population is available to encode the scene. These recipient neurons can act as coincidence detectors for spikes impinging from different afferent fibers (Alonso et al., 1996; Usrey et al., 2000). Thus, the early visual cortex may well decode the distrib-

uted firing patterns among optic nerve fibers and represent them explicitly again in a single-neuron code.

One cost associated with this coding scheme are the additional action potentials needed to transmit the synchronous firing patterns. For example, an interneuron connected to three ganglion cells produces three spikes, whereas only one would be needed if the interneuron had its own optic nerve fiber. The redundancy introduced by the extra spikes is precisely what allowed our algorithm to find them in the first place. It has been argued that the metabolic energy required to maintain spiking is a serious constraint on retinal function (Balasubramanian et al., 2001; Laughlin, 2001). The predominance of synchronous firing suggests that this is not the only constraint, because many action potentials could be saved with a “labeled line” coding scheme. Instead, the system finds itself under competing pressures, among them the need for acute vision, the need for eye movements to stabilize gaze or to scan the scene, the consequent mechanical constraints on the thickness of the optic nerve, and the cost of metabolic energy to sustain retinal function.

Experimental Procedures

Multielectrode Recording and Visual Stimulation

The preparation of the retina, visual stimulation, and recording of ganglion cell spike trains were conducted as described (Meister et al., 1994; Smimakis et al., 1997). In brief, a piece of isolated retina was placed ganglion cell side down onto a flat array of 61 extracellular electrodes spaced at 60 μm distances. Spikes from the ganglion cell layer were recorded on all channels in parallel and sorted into single-unit spike trains. The spatial location of a ganglion cell body was triangulated from the electrodes that recorded its spikes, with each electrode location weighted by the corresponding spike amplitude.

Visual stimuli were generated on a computer monitor and projected demagnified onto the retina. For random checkerboard stimulation, the field was divided into a grid of 70 μm squares. Within each square, the red, green, and blue guns of the monitor were turned on or off by a random choice. New random values were chosen periodically at time intervals of 15–120 ms. The mean light intensity of the stimulus ranged from high scotopic to low photopic.

Identification of Synchronous Spiking Groups

We identified groups of synchronized cells in the recorded spike trains by seeking a so-called factorial recoding of the data that would compress the spike trains. This is a common approach for finding and removing patterns in a data set for the purpose of compression (Storer, 1988). Specifically, we adapted a method due to Redlich (1993) that is suited for data containing strong correlations within certain subsets of the symbol set but few correlations between these subsets. In multineuron spike trains, such strongly correlated subsets correspond to cells that fire together as a group.

The ganglion cell spike trains were binned into time intervals of width 50 ms, and represented in a binary fashion:

$$r^{(A)}_j = \begin{cases} 1, & \text{if cell } A \text{ fired } \geq 1 \text{ spikes in time bin } j \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

If two cells A and B fire together frequently, they will tend to each exhibit 1 within the same time bin. This representation is inefficient in that two 1's are used to indicate the occurrence of a single event. One can recode the data by defining a new symbolic cell, AB , whose spikes represent synchronous firing of the real cells A and B (Figure 2B).

$$r^{(AB)}_j = r^{(A)}_j r^{(B)}_j = \begin{cases} 1, & \text{if cells } A \text{ and } B \text{ fired in time bin } j \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

The events where either A or B fires alone are retained in the symbols $A'B$ and AB' .

This new representation uses fewer 1's to specify the same data, but also requires an additional spike train. To see whether this recoding really enables a compression, we compute the bit entropy of the data set before and after the symbol change. This corresponds to the minimal number of bits per time bin required to store all the individual spike trains (Shannon and Weaver, 1963):

$$H = -\sum_{i=1}^M P_i \log P_i + (1 - P_i) \log(1 - P_i), \quad (3)$$

where M is the number of spike trains, P_i is the probability that symbol i has a spike in an arbitrary time bin,

$$P_i = \frac{1}{N} \sum_{j=1}^N r_j^{(i)}, \quad (4)$$

N is the number of time bins in the data set, and \log denotes the logarithm to base 2.

Introducing a new symbol AB adds a term to the sum of Equation 3, but removing the synchronous spikes from the remaining spike trains $A'B$ and AB' decreases their contributions to the sum. Denoting the probability that A and B fire together by

$$P_{AB} = \frac{1}{N} \sum_{j=1}^N r_j^{(AB)}, \quad (5)$$

one finds that the net reduction in entropy from the recoding is

$$\begin{aligned} \Delta H_{AB} = & \log \left(\frac{(1 - P_{AB})(1 - P_A + P_{AB})(1 - P_B + P_{AB})}{(1 - P_A)(1 - P_B)} \right) \\ & + P_A \log \left(\frac{(P_A - P_{AB})(1 - P_A)}{(1 - P_A + P_{AB})P_A} \right) \\ & + P_B \log \left(\frac{(P_B - P_{AB})(1 - P_B)}{(1 - P_B + P_{AB})P_B} \right) \\ & + P_{AB} \log \left(\frac{P_{AB}(1 - P_B + P_{AB})(1 - P_A + P_{AB})}{(1 - P_{AB})(P_A - P_{AB})(P_B - P_{AB})} \right) \end{aligned} \quad (6)$$

If ΔH_{AB} is positive, one can compress the data set by representing the joint firing of A and B explicitly in AB , and the value of ΔH_{AB} is the number of data storage bits per time bin one saves by making the replacement.

To gain some intuition for this criterion, consider the case of $P_{AB} \ll P_A, P_B$, where only a small fraction of each cell's spikes are joint firings (Redlich, 1993). Under those conditions, Equation 6 simplifies to

$$\Delta H_{AB} \approx P_{AB} \log(P_{AB}/P_A P_B). \quad (7)$$

In this limit, one sees that the benefits of defining the new symbol AB are greater the more often the joint firing event occurs, due to the factor P_{AB} . Additionally, the factor $\log(P_{AB}/P_A P_B)$ is large when joint firing occurs much more often than predicted from the product of the individual firing rates. Thus, our measure of a firing pattern's importance is a combination of its absolute frequency and its relative unexpectedness given the firing rates of the constituent neurons.

The identification of groups begins by computing ΔH_{AB} for all cell pairs. The pair yielding the largest value is chosen and its synchronous spikes are given a new symbol AB , as described above. Then this process is repeated. In each round, the symbolic cells of previous iterations are treated on an equal footing with the real cells. After just two iterations, it is possible to combine two symbolic cells into a new one and thus identify synchronized groups of three or more real cells. The iterations stop when the largest available ΔH_{AB} falls below a predetermined threshold of significance.

We set this threshold high enough so the algorithm would not identify groups based on chance coincidences. To find this threshold, each ganglion cell's spike train was shifted by a different random number of time bins. This shuffling procedure maintains the actual

interspike interval distribution and all other single-cell statistics of the spike trains (Oram et al., 1999), but reduces the synchronous spiking events between different neurons to the spurious level expected by chance. We computed the largest ΔH_{AB} in this shuffled data set and chose a threshold just above this critical value. This ensured that no artifactual groups appeared when the algorithm processed the original set of spike trains.

To assess the reliability of the group identification further, we occasionally repeated the analysis using only half of the data set. About 80% of the patterns identified were identical between the two analyses; small discrepancies arose because the accretion of cells into groups (Figure 2B) proceeded in a slightly different order depending on the joint frequencies in the two data sets. The overall set of patterns emerging from the two analyses had very similar statistics. For example, the number of partners engaged by a ganglion cell (Figure 4C) differed by 1 or less for 90% of the neurons. We also tested the dependence on the width of the coincidence time bin, with values down to 20 ms. This again yielded very similar results, verifying once more that the identified patterns are not spurious.

Strength of the Correlation in Groups

To compute the correlation index (Figure 3) among M cells labeled $1 \dots M$, note that their frequency of synchronous firing is (Equation 1)

$$P_{1\dots M} = \frac{1}{N} \sum_{j=1}^N \prod_{i=1}^M r_i^{j_i}, \quad (8)$$

whereas the frequency expected if they were firing independently is

$$P_{1\dots M} = \prod_{i=1}^M \frac{1}{N} \sum_{j=1}^N r_i^{j_i}. \quad (9)$$

For each group we computed the ratio of these two terms, the correlation index

$$C_{1\dots M} = P_{1\dots M} / \prod_{i=1}^M P_i. \quad (10)$$

For Figure 3, inset, we computed $P_{ABC}/(P_{AB}P_C)$, the frequency of a triplet ABC compared to that of the single cell C and the pair AB from which the triplet was identified by the search algorithm.

Time Binning

The 50 ms time bin used in our pattern search captures the narrow and intermediate correlations found in the salamander retina (Figure 1B; Brivanlou et al., 1998) and is less sensitive to the broad correlations. Note that if two cells fire 25 ms apart, they will be detected in the same bin half the time.

Once the identity of the various cell groups was known, we returned to the original spike trains to parse them into events corresponding to the firing of these groups. The purpose was to remove the time bin boundaries introduced in Equation 1 by sliding a time window continuously over the spike trains and testing for occurrence of group firing within the window. One cell in each group was denoted the reference cell. If spikes from all other cells in the group fell within ± 25 ms of a spike from the reference cell, that set of spikes was denoted as a firing of the group. After scanning the spike trains this way for each group, some spikes remained unassigned. This procedure led to the fraction of spikes in groups reported in Figure 4.

Estimating the Effects of Undersampling

Because the electrode array does not record from all the overlying cells, the observed groups of synchronized cells will tend to be smaller than the actual groups. Suppose the recording method is unbiased, such that each ganglion cell has a probability f of being recorded. For a cell group of actual size L , the probability that it appears in the recording with size M is specified by the binomial distribution:

$$p(M|L) = \frac{L!}{(L-M)!M!} f^M (1-f)^{L-M}. \quad (11)$$

Denoting the fraction of groups with actual size L by $p(L)$, the fraction of groups with observed size M is

$$\begin{aligned} p(M) &= \sum_{L=M}^{\infty} p(L)p(M|L) \\ &= \sum_{L=M}^{\infty} p(L) \frac{L!}{(L-M)!M!} f^M (1-f)^{L-M}. \end{aligned} \quad (12)$$

With this, the true distribution of group size $p(L)$ can be found from the observed distribution $p(M)$ recursively, starting with the groups of the largest observed size. For our observed distribution (Figure 4A), one obtains reasonable actual distributions, with positive numbers for $p(L)$, only for large values of the recording probability ($f \geq 0.75$ in darkness or $f \geq 0.5$ under visual stimulation). Because the overall recording yield is likely much smaller ($f \leq 0.2$) (Meister et al., 1994), this calculation suggests that the assumption of unbiased cell sampling is incorrect. Synchronized firing tends to involve ganglion cells of the same functional type (Mastronarde, 1989; Meister et al., 1995; DeVries, 1999), and the electrode array may record some of these types with a high efficiency of $f \geq 0.5$. Other cell types must be recorded with lower efficiency, and correspondingly one learns less about their synchronous firing patterns.

Computing the STA and Receptive Field Parameters

For analyzing the flickering checkerboard experiments, it is useful to normalize the light stimulus by subtracting the mean and dividing by the standard deviation,

$$s = \text{normalized light intensity} = \frac{I - \langle I \rangle}{\langle (I^2) \rangle - \langle I \rangle^2}^{1/2}. \quad (13)$$

If the values of the random checkerboard stimulus are

$$s_{ijk} = \text{stimulus value for time bin } i, \text{ region } j, \text{ and monitor gun } k, \quad (14)$$

and the resulting response of a retinal ganglion cell is

$$r_i = \text{number of spikes fired in time bin } i, \quad (15)$$

then the spike-triggered average stimulus is computed as

$$\begin{aligned} h_{ijk} &= \text{spike-triggered average stimulus in time bin } i \\ &\text{relative to the spike, region } j, \text{ and monitor gun } k. \\ &= \frac{\sum_{i=1}^N r_i s_{i+jk}}{\sum_{i=1}^N r_i} \end{aligned} \quad (16)$$

This STA is a function of time before the spike, spatial location, and color channel. To obtain an estimate of the spatial receptive field, we approximated the STA by a product of three functions that individually depend only on time, space, and the color channel:

$$\begin{aligned} h_{ijk} &\approx a_i b_j c_k \\ a_i &= a(t_i) = \text{time course of the STA in time bin } i \\ b_j &= b(x_j) = \text{profile of the STA in region } j \\ c_k &= \text{sensitivity of the STA to color channel } k. \end{aligned} \quad (17)$$

This was done by finding the three functions a_i , b_j , and c_k that minimize the squared error

$$\sum_{ijk} (h_{ijk} - a_i b_j c_k)^2.$$

Although the responses of retinal ganglion cells do not exactly separate in space, time, and color (Wandell, 1995; Meister and Berry, 1999), this approximation provided a good estimate of the spatial receptive field. The shape of this profile $b(x_j)$ was then fit with a two-dimensional Gaussian bell shape,

$$b(x) = \text{Bexp}\left(-\frac{1}{2}(\mathbf{x} - \mathbf{u})^T \mathbf{C}^{-1}(\mathbf{x} - \mathbf{u})\right), \quad (18)$$

where

$$\begin{aligned} \mathbf{u} &= \text{center of the Gaussian} \\ \mathbf{C} &= \text{covariance matrix for the shape of the Gaussian.} \end{aligned} \quad (19)$$

From this, we defined the radius of the receptive field as the mean

radius of the Gaussian at 1 standard deviation from the center

$$R = \text{mean radius of the Gaussian} = (\det \mathbf{C})^{1/4}. \quad (20)$$

Variation of the STA in a Group of Cells

The variation among the STAs in a given group of cells was computed separately for the time course a_i , the spatial profile b_j , and the spectral sensitivity c_k . For example, let

$$\mathbf{a} = [\mathbf{a}_i] / \sqrt{\sum_i \mathbf{a}_i^2}$$

be the normalized vector specifying the shape of a given cell's time course. For any group of M cells, we measured the scatter of their sensitivity vectors about the mean and divided by the length of the mean to obtain the variation V :

$$V = \sqrt{\frac{1}{M-1} \sum_{i=1}^M (\mathbf{a}^{(i)} - \mathbf{m})^2 / \mathbf{m}^2}, \text{ where } \mathbf{m} = \frac{1}{M} \sum_{i=1}^M \mathbf{a}^{(i)}. \quad (21)$$

For the spatial profile and the spectral sensitivity, the same analysis was performed on the vectors b_j and c_k , respectively.

The STA for an LN Neuron

To develop an expectation for the STA of a multicell spike pattern, it helps to work through a specific model of the neuron's light response. Here we assume that the ganglion cell firing rate $r(t)$ depends on the stimulus $s(t)$ through a so-called "linear-nonlinear" (LN) relationship (Chichilnisky, 2001).

$$r(t) = \text{firing rate at time } t = N\left(s(t + t', \mathbf{x}, \lambda) f(t', \mathbf{x}, \lambda) dt'\right)$$

where

$s(t, \mathbf{x}, \lambda)$ = stimulus at time t , location \mathbf{x} , wavelength λ

$f(t, \mathbf{x}, \lambda)$ = spatial-temporal-spectral filter

$N(\)$ = instantaneous nonlinear function (22)

In the above discrete notation,

$$r_l = r(t_l) = \text{firing rate in time bin } l = N\left(\sum_{i,j,k} s_{l+i,j,k} f_{i,j,k}\right). \quad (23)$$

At any instant in time, the recent stimulus values are averaged over time, space, and color channels, with a linear weighting function given by f_{ijk} . The result is transformed by an instantaneous nonlinear function $N(\)$ to yield the firing rate.

One also needs to spell out the statistics of the stimulus: here we assume for simplicity that the stimulus values for the flickering checkerboard are drawn from a normal Gaussian distribution, so the probability of getting any given stimulus value is

$$p(s_{ijk}) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} s_{ijk}^2\right), \quad (24)$$

where the different s_{ijk} are drawn independently of each other.

What is the STA of a neuron that behaves according to Equation 23 under the stimulus ensemble of Equation 24? At any given point in time, say $l = 0$ in Equation 23, the firing rate depends on the stimulus values in the preceding time interval of perhaps 1 s length. One can think of these stimulus values s_{ijk} as a large-dimensional vector $\mathbf{s} = [s_{ijk}]$. The probability distribution of this stimulus vector is a multidimensional Gaussian,

$$p(\mathbf{s}) \propto \exp\left(-\frac{1}{2} \mathbf{s}^2\right), \quad (25)$$

and for any given \mathbf{s} , the instantaneous firing rate is

$$r(\mathbf{s}) = N(\mathbf{s} \cdot \mathbf{f}), \quad (26)$$

where $\mathbf{f} = [f_{ijk}]$ is the vector representation of the neuron's linear filter. The spike-triggered average of the stimulus vector is obtained by weighting its distribution with the firing rate of the neuron,

$$\begin{aligned} \mathbf{h} &= \text{spike-triggered average stimulus vector} \\ &= \frac{\int \mathbf{s} p(\mathbf{s}) r(\mathbf{s}) d^n \mathbf{s}}{\int p(\mathbf{s}) r(\mathbf{s}) d^n \mathbf{s}} = \frac{\int \mathbf{s} \exp\left(-\frac{1}{2} \mathbf{s}^2\right) N(\mathbf{s} \cdot \mathbf{f}) d^n \mathbf{s}}{\int \exp\left(-\frac{1}{2} \mathbf{s}^2\right) N(\mathbf{s} \cdot \mathbf{f}) d^n \mathbf{s}} \end{aligned} \quad (27)$$

On the left-hand side, \mathbf{h} is a vector. On the right-hand side, the distribution $p(\mathbf{s})$ is spherically symmetric, but \mathbf{f} is a vector. From symmetry arguments one concludes immediately that \mathbf{h} points in the direction of \mathbf{f} , so the neuron's spike-triggered average is proportional to its filter function,

$$\mathbf{h}_{ijk} = \alpha f_{ijk}. \quad (28)$$

The specific value of α depends on the function $N(\)$.

The STA for a Synchronous Firing Pattern among Independent LN Neurons

Suppose two retinal ganglion cells A and B each respond to the stimulus according to an LN model (Figure 8C). Their instantaneous firing rates are

$$r^{(A)}(\mathbf{s}) = N^{(A)}(\mathbf{s} \cdot \mathbf{f}^{(A)}) \text{ and } r^{(B)}(\mathbf{s}) = N^{(B)}(\mathbf{s} \cdot \mathbf{f}^{(B)}), \quad (29)$$

where $\mathbf{f}^{(A)}$ and $\mathbf{f}^{(B)}$ are the respective linear filters, and $N^{(A)}(\)$ and $N^{(B)}(\)$ are the nonlinearities. If the two cells operate independently, then the probability that they fire synchronously is simply proportional to the product of their individual firing rates. Specifically, if spikes are considered synchronous within a short time interval Δt , then the rate of occurrence of synchronous pairs is

$$r^{(AB)}(\mathbf{s}) = \Delta t r^{(A)}(\mathbf{s}) r^{(B)}(\mathbf{s}) = \Delta t N^{(A)}(\mathbf{s} \cdot \mathbf{f}^{(A)}) N^{(B)}(\mathbf{s} \cdot \mathbf{f}^{(B)}). \quad (30)$$

Therefore, the STA of the synchronous pair is found analogously to Equation 27

$$\begin{aligned} \mathbf{h}^{(AB)} &= \frac{\int \mathbf{s} p(\mathbf{s}) r^{(AB)}(\mathbf{s}) d^n \mathbf{s}}{\int p(\mathbf{s}) r^{(AB)}(\mathbf{s}) d^n \mathbf{s}} \\ &= \frac{\int \mathbf{s} \exp\left(-\frac{1}{2} \mathbf{s}^2\right) N^{(A)}(\mathbf{s} \cdot \mathbf{f}^{(A)}) N^{(B)}(\mathbf{s} \cdot \mathbf{f}^{(B)}) d^n \mathbf{s}}{\int \exp\left(-\frac{1}{2} \mathbf{s}^2\right) N^{(A)}(\mathbf{s} \cdot \mathbf{f}^{(A)}) N^{(B)}(\mathbf{s} \cdot \mathbf{f}^{(B)}) d^n \mathbf{s}} \end{aligned} \quad (31)$$

Note that the only vector quantities on the right hand side are $\mathbf{f}^{(A)}$ and $\mathbf{f}^{(B)}$, which by Equation 28 are proportional to $\mathbf{h}^{(A)}$ and $\mathbf{h}^{(B)}$, respectively. The vector $\mathbf{h}^{(AB)}$ must be a linear combination of these two, and therefore the STA of a synchronous spike pair is a weighted average of the STAs from the individual cells,

$$\mathbf{h}^{(AB)} = \alpha \mathbf{h}^{(A)} + \beta \mathbf{h}^{(B)}, \quad (32)$$

where α and β specify the relative weightings. If the two retinal ganglion cells are of the same functional type and thus use the same nonlinear function $N(\)$, then symmetry dictates that $\alpha = \beta$, and the STA of the pair is

$$\mathbf{h}^{(AB)} = \alpha (\mathbf{h}^{(A)} + \mathbf{h}^{(B)}). \quad (33)$$

These arguments extend immediately to groups of three and more neurons. For neurons that respond independently according to the LN model, the STA of a synchronous firing pattern is an average, possibly weighted, of the STAs of the component neurons.

We comment briefly on this result, because it conflicts with a widely held expectation, "If two neurons respond independently of each other, then the only way to get responses from both is to flash light into both receptive fields. So shouldn't the receptive field of joint spikes be the overlap region of the receptive fields?" This natural intuition derives from the classical method of probing the receptive fields of visual neurons, namely with a single small, flashing spot moved to different locations. Under those conditions, the only way to get both neurons to respond is indeed to place the spot in the overlap region. When one analyzes the receptive fields of joint spikes, one obtains by necessity the overlap region of the individual receptive fields (Ghose et al., 1994).

In the present work, the receptive fields were mapped with a white-noise flicker stimulus, in which spots are flashing continuously, randomly, and independently at every location in the visual field. Two neurons can now fire synchronously either because a spot flashed in the overlap region or because two spots flashed in the disjoint regions. Even two neurons with no receptive field overlap will fire together when their receptive fields are stimulated at the same time; thus, the receptive field of their joint spikes is the sum of their individual receptive fields. The proof leading to Equation 33 shows that this relationship holds even when the individual receptive

fields overlap, assuming an LN model for the individual neurons. The difference from the expectation discussed above arises entirely from the use of a different stimulus ensemble. Receptive field mapping by random flicker stimulation is becoming a standard method in visual neuroscience; in interpreting the results, it is important to consider how they depend on the stimulus statistics.

The STA for Neurons that Share a Presynaptic Spike Train

Suppose two retinal ganglion cells *A* and *B* share excitation from a third presynaptic neuron *C*, such that when the presynaptic cell fires, both postsynaptic ganglion cells fire together (Figure 8B). In addition, *A* receives unshared input from *D*, and *B* receives unshared input from *E*. For concreteness, suppose that the interneurons *C*, *D*, and *E* all respond to the stimulus by an LN mechanism.

The firing rate of a ganglion cell is simply the sum of the rate of its inputs,

$$\begin{aligned} r^{(A)}(\mathbf{s}) &= r^{(C)}(\mathbf{s}) + r^{(D)}(\mathbf{s}) \\ &= N^{(C)}(\mathbf{s} \cdot \mathbf{f}^{(C)}) + N^{(D)}(\mathbf{s} \cdot \mathbf{f}^{(D)}) \\ r^{(B)}(\mathbf{s}) &= r^{(C)}(\mathbf{s}) + r^{(E)}(\mathbf{s}) \\ &= N^{(C)}(\mathbf{s} \cdot \mathbf{f}^{(C)}) + N^{(E)}(\mathbf{s} \cdot \mathbf{f}^{(E)}), \end{aligned} \quad (34)$$

where we assume that there is no temporal integration within the ganglion cell. In practice, the arguments below apply if the ganglion cell's integration time for these spiking inputs is less than the correlation interval of 50 ms. Indeed, Kim and Rieke (2001) measured the integration time for membrane currents at ~ 15 ms.

By the symmetry arguments used above, one finds that the individual STAs of cells *A* and *B* are linear averages of those of the presynaptic neurons,

$$\begin{aligned} \mathbf{h}^{(A)} &= \alpha^{(C)}\mathbf{h}^{(C)} + \alpha^{(D)}\mathbf{h}^{(D)} \\ \mathbf{h}^{(B)} &= \beta^{(C)}\mathbf{h}^{(C)} + \beta^{(E)}\mathbf{h}^{(E)}. \end{aligned} \quad (35)$$

If the interneurons *C*, *D*, and *E* are of the same functional type and use the same nonlinear function $N(\cdot)$, then the additional symmetry provides

$$\begin{aligned} \mathbf{h}^{(A)} &= (\mathbf{h}^{(C)} + \mathbf{h}^{(D)})/2 \\ \mathbf{h}^{(B)} &= (\mathbf{h}^{(C)} + \mathbf{h}^{(E)})/2. \end{aligned} \quad (36)$$

The two ganglion cells *A* and *B* fire synchronously whenever their shared input *C* fires. If the coincidence interval Δt is short compared to the average interval between spikes, accidental coincidences from unshared inputs *D* and *E* are rare, and

$$r^{(AB)}(\mathbf{s}) = r^{(C)}(\mathbf{s}), \quad (37)$$

implying

$$\mathbf{h}^{(AB)} = \mathbf{h}^{(C)}. \quad (38)$$

In practice, the fraction of coincidences among the interneurons *C*, *D*, and *E* that are accidental will be comparable to the fraction among different ganglion cells, namely the inverse of the correlation index (Figure 3). This is $<10\%$ for half the spike pairs, and $<1\%$ for half the spike triplets; thus, the approximation leading to Equation 37 holds to good accuracy.

Comparing Equations 36 and 38, one sees that the receptive field of synchronous spikes $\mathbf{h}^{(AB)}$ is a subregion of the individual receptive fields, $\mathbf{h}^{(A)}$ and $\mathbf{h}^{(B)}$, of each of the component neurons. This contrasts with the result for spike coincidences among neurons that operate independently (Equation 33), whose receptive field is broader than those of the individual cells. Again, the treatment extends readily to larger groups of synchronized cells.

Acknowledgments

This work was supported by NIH grant EY10020 (M.M.) and by Bell Laboratories (M.J.S.). We thank the members of our group for comments on the manuscript.

Received: March 12, 2002
Revised: November 21, 2002

References

- Abeles, M., and Gerstein, G.L. (1988). Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *J. Neurophysiol.* 60, 909–924.
- Alonso, J.M., Usrey, W.M., and Reid, R.C. (1996). Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature* 383, 815–819.
- Arnett, D.W. (1978). Statistical dependence between neighboring retinal ganglion cells in goldfish. *Exp. Brain Res.* 32, 49–53.
- Arnett, D., and Spraker, T.E. (1981). Cross-correlation analysis of the maintained discharge of rabbit retinal ganglion cells. *J. Physiol.* 317, 29–47.
- Balasubramanian, V., Kimber, D., and Berry, M.J. (2001). Metabolically efficient information processing. *Neural Comput.* 13, 799–815.
- Brivanlou, I.H., Warland, D.K., and Meister, M. (1998). Mechanisms of concerted firing among retinal ganglion cells. *Neuron* 20, 527–539.
- Buzsaki, G., Bickford, R.G., Ryan, L.J., Young, S., Prohaska, O., Mandel, R.J., and Gage, F.H. (1989). Multisite recording of brain field potentials and unit activity in freely moving rats. *J. Neurosci. Methods* 28, 209–217.
- Chapin, J.K., and Nicolelis, M.A. (1999). Principal component analysis of neuronal ensemble activity reveals multidimensional somatosensory representations. *J. Neurosci. Methods* 94, 121–140.
- Chichilnisky, E.J. (2001). A simple white noise analysis of neuronal light responses. *Network* 12, 199–213.
- Cleland, B.G., Dubin, M.W., and Levick, W.R. (1971). Sustained and transient neurones in the cat's retina and lateral geniculate nucleus. *J. Physiol.* 217, 473–496.
- Cleland, B.G., Levick, W.R., Morstyn, R., and Wagner, H.G. (1976). Lateral geniculate relay of slowly conducting retinal afferents to cat visual cortex. *J. Physiol.* 255, 299–320.
- Dan, Y., Alonso, J.M., Usrey, W.M., and Reid, R.C. (1998). Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nat. Neurosci.* 1, 501–507.
- Della Santina, C.C., Kovacs, G.T., and Lewis, E.R. (1997). Multi-unit recording from regenerated bullfrog eighth nerve using implantable silicon-substrate microelectrodes. *J. Neurosci. Methods* 72, 71–86.
- DeVries, S. (1999). Correlated firing in rabbit retinal ganglion cells. *J. Neurophysiol.* 81, 908–920.
- Frostig, R.D., Frostig, Z., and Harper, R.M. (1990). Recurring discharge patterns in multiple spike trains. I. Detection. *Biol. Cybern.* 62, 487–493.
- Gerstein, G.L., and Aertsen, A.M. (1985). Representation of cooperative firing activity among simultaneously recorded neurons. *J. Neurophysiol.* 54, 1513–1528.
- Gerstein, G.L., Perkel, D.H., and Subramanian, K.N. (1978). Identification of functionally related neural assemblies. *Brain Res.* 140, 43–62.
- Ghose, G.M., Ohzawa, I., and Freeman, R.D. (1994). Receptive-field maps of correlated discharge between pairs of neurons in the cat's visual cortex. *J. Neurophysiol.* 71, 330–346.
- Hunter, I.W., and Korenberg, M.J. (1986). The identification of nonlinear biological systems: Wiener and Hammerstein cascade models. *Biol. Cybern.* 55, 135–144.
- Johnsen, J.A., and Levine, M.W. (1983). Correlation of activity in neighbouring goldfish ganglion cells: relationship between latency and lag. *J. Physiol.* 345, 439–449.
- Kim, K.J., and Rieke, F. (2001). Temporal contrast adaptation in the input and output signals of salamander retinal ganglion cells. *J. Neurosci.* 21, 287–299.
- Kovacs, G.T., Stormont, C.W., Halks-Miller, M., Belczynski, C.R., Jr., Della Santina, C.C., Lewis, E.R., and Maluf, N.I. (1994). Silicon-substrate microelectrode arrays for parallel recording of neural activity in peripheral and cranial nerves. *IEEE Trans. Biomed. Eng.* 41, 567–577.

- Laughlin, S.B. (2001). Energy as a constraint on the coding and processing of sensory information. *Curr. Opin. Neurobiol.* *11*, 475–480.
- Mastrorarde, D.N. (1983). Correlated firing of cat retinal ganglion cells. I. Spontaneously active inputs to X- and Y-cells. *J. Neurophysiol.* *49*, 303–324.
- Mastrorarde, D.N. (1989). Correlated firing of retinal ganglion cells. *Trends Neurosci.* *12*, 75–80.
- Meister, M. (1996). Multineuronal codes in retinal signaling. *Proc. Natl. Acad. Sci. USA* *93*, 609–614.
- Meister, M., 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.
- Meister, M., Lagnado, L., and Baylor, D.A. (1995). Concerted signaling by retinal ganglion cells. *Science* *270*, 1207–1210.
- Nordhausen, C.T., Maynard, E.M., and Normann, R.A. (1996). Single unit recording capabilities of a 100 microelectrode array. *Brain Res.* *726*, 129–140.
- Oram, M.W., Wiener, M.C., Lestienne, R., and Richmond, B.J. (1999). Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *J. Neurophysiol.* *81*, 3021–3033.
- Redlich, A.N. (1993). Redundancy reduction as a strategy for unsupervised learning. *Neural Comput.* *5*, 289–304.
- Shannon, C.E., and Weaver, W. (1963). *The Mathematical Theory of Communication*, 2nd edition (Chicago: University of Illinois Press).
- 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.
- Storer, J.A. (1988). *Data Compression: Methods and Theory* (Rockville, MD: Computer Science Press).
- Usrey, W.M., and Reid, R.C. (1999). Synchronous activity in the visual system. *Annu. Rev. Physiol.* *61*, 435–456.
- Usrey, W.M., Alonso, J.M., and Reid, R.C. (2000). Synaptic interactions between thalamic inputs to simple cells in cat visual cortex. *J. Neurosci.* *20*, 5461–5467.
- Wandell, B.A. (1995). *Foundations of Vision* (Sunderland, MA: Sinauer).
- Warland, D.K., Reinagel, P., and Meister, M. (1997). Decoding visual information from a population of retinal ganglion cells. *J. Neurophysiol.* *78*, 2336–2350.
- Wilson, M.A., and McNaughton, B.L. (1993). Dynamics of the hippocampal ensemble code for space. *Science* *261*, 1055–1058.