

The Structure of Multi-Neuron Firing Patterns in Primate Retina

Jonathon Shlens,¹ Greg D. Field,¹ Jeffrey L. Gauthier,¹ Matthew I. Grivich,² Dumitru Petrusca,² Alexander Sher,² Alan M. Litke,² and E. J. Chichilnisky¹

¹Department of Systems Neurobiology, The Salk Institute, La Jolla, California 92037, and ²Santa Cruz Institute for Particle Physics, University of California, Santa Cruz, California 95064

Current understanding of many neural circuits is limited by our ability to explore the vast number of potential interactions between different cells. We present a new approach that dramatically reduces the complexity of this problem. Large-scale multi-electrode recordings were used to measure electrical activity in nearly complete, regularly spaced mosaics of several hundred ON and OFF parasol retinal ganglion cells in macaque monkey retina. Parasol cells exhibited substantial pairwise correlations, as has been observed in other species, indicating functional connectivity. However, pairwise measurements alone are insufficient to determine the prevalence of multi-neuron firing patterns, which would be predicted from widely diverging common inputs and have been hypothesized to convey distinct visual messages to the brain. The number of possible multi-neuron firing patterns is far too large to study exhaustively, but this problem may be circumvented if two simple rules of connectivity can be established: (1) multi-cell firing patterns arise from multiple pairwise interactions, and (2) interactions are limited to adjacent cells in the mosaic. Using maximum entropy methods from statistical mechanics, we show that pairwise and adjacent interactions accurately accounted for the structure and prevalence of multi-neuron firing patterns, explaining $\sim 98\%$ of the departures from statistical independence in parasol cells and $\sim 99\%$ of the departures that were reproducible in repeated measurements. This approach provides a way to define limits on the complexity of network interactions and thus may be relevant for probing the function of many neural circuits.

Key words: vision; information theory; correlated variability; neural coding; synchrony; retinal ganglion cell

Introduction

A central challenge in neuroscience is to understand how large circuits of neurons represent and process information. For decades, studies of neural function were restricted to recordings from single neurons, with the tacit assumption that the function of complex circuits could be deciphered with such measurements (Wandell, 1995). However, multi-neuron recordings have revealed substantial interactions that cannot be observed with single-neuron recording. For example, retinal ganglion cells (RGCs) exhibit strong stimulus-independent correlated activity; the circuits mediating such correlations and the consequences for visual processing are not fully understood (Arnett, 1978; Johnsen and Levine, 1983; Mastronarde, 1983b; Meister et al., 1995; DeVries and Baylor, 1997; Nirenberg et al., 2001; Schneidman et al., 2003a). Such findings suggest interesting possibilities for circuit function but at the same time raise a major concern: will it be necessary to record from all of the cells in a neural circuit, and analyze all possible interactions, to determine how the circuit works? If so, a deep understanding of many neural systems may be out of reach for a long time.

In this context, any simplifying principles that can make the problem more tractable are of great value. To illustrate this, consider the number of possible circuits terminating on a collection of n neurons. The number of distinct circuits is determined by the number of distinct input patterns that contribute to the circuit. Examples of distinct input patterns are illustrated schematically for $n = 8$ cells in Figure 1A. Even with this small value of n , the entire collection of distinct patterns is too large to depict easily; thus, only selected examples are shown. In general, the number of possible input patterns is $\sim 2^n$, a prohibitive complexity: for the collection of several hundred RGCs depicted in Figure 2A, 2^n exceeds the number of stars in the known universe (Liske et al., 2003). However, two simple constraints on connectivity can dramatically simplify the problem. The first is “pairwise” connectivity, in which each input pattern contacts only two cells. All pairwise patterns may be depicted together in a single diagram (see Fig. 1B). The second is “adjacent” connectivity, in which each pattern contacts all cells within its extent. This constraint also restricts the possibilities to a set that can be depicted easily (see Fig. 1C). The pairwise and adjacent constraints each reduce the number of possible patterns to $\sim n^2$, a huge simplification. With both constraints, the number of possible patterns is reduced to

Received March 24, 2006; revised June 27, 2006; accepted July 1, 2006.

This work was supported by National Science Foundation (NSF) Integrative Graduate Education and Research Traineeship Grant DGE-0333451, a Burroughs-Wellcome La Jolla Interfaces in Science Fellowship (J.S.), NSF Grant PHY-0417175 (A.M.L.), a Sloan Research Fellowship, and National Institutes of Health Grant EY13150 (E.J.C.). We thank E. Simoncelli for key discussions; W. Dabrowski, A. Grillo, P. Grybos, P. Hottowy, and S. Kachiguine for technical development; R. Malouf for valuable technical suggestions; P. Latham for comments on this manuscript; E. Callaway, H. Fox, and K. Osborn for providing access to retinas; B. Kutka for technical assistance; and S. Barry for machining. We thank the San Diego Supercomputer Center and the National Science Foundation (Cooperative Agreements 05253071 and 0438741) for large-scale data storage.

Correspondence should be addressed to Jonathon Shlens, Systems Neurobiology, The Salk Institute, 10010 North Torrey Pines Road, La Jolla, CA 92037. E-mail: shlens@salk.edu.

DOI:10.1523/JNEUROSCI.1282-06.2006

Copyright © 2006 Society for Neuroscience 0270-6474/06/268254-13\$15.00/0

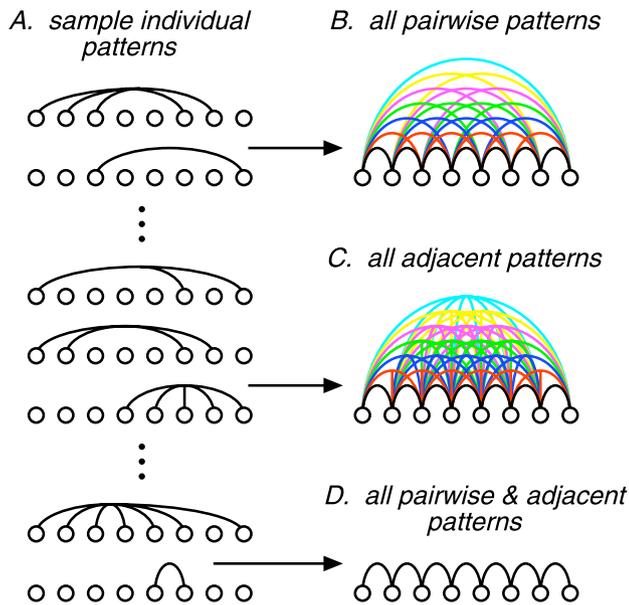


Figure 1. Patterns of connectivity. **A**, A small random sample of the possible input patterns to a collection of $n = 8$ cells. From top to bottom, seven different hypothetical input patterns are shown, each terminating on a different collection of target cells (circles). In general, the total number of distinct patterns is on the order of 2^n . **B**, All distinct pairwise patterns, superimposed. Each color shows, superimposed, all pairwise patterns across cells separated by a particular distance. An example of a single pairwise pattern from **A** is indicated with an arrow. In general, the number of distinct pairwise patterns is on the order of n^2 . In this simplified diagram, the 28 patterns shown exclude wraparound and symmetric patterns. **C**, All distinct adjacent patterns, superimposed. Each color shows, superimposed, all adjacent patterns consisting of a particular number of cells. An example of one adjacent pattern from **A** is indicated with an arrow. Again, 28 distinct patterns are shown. In general, the number of distinct patterns is on the order of n^2 . **D**, All possible patterns that are both pairwise and adjacent, superimposed. Seven patterns are shown, and a single example is indicated with an arrow. In general, the number of possible patterns is on the order of n .

$\sim n$ (see Fig. 1D). The practical consequences of this simplification are profound: in principle, one can understand the function of the entire circuit simply by recording from individual cells and pairs of neighboring cells.

The importance of these simplifying principles is illustrated in the retina. Recent work suggests that synchronized firing among RGCs could originate in common input to multiple RGCs, forming a multiplexed neural code in which the large and distinctive electrical “footprint” of a presynaptic cell conveys a specific visual message to the brain (Meister et al., 1995; Schnitzer and Meister, 2003). Other studies have also suggested interactions in RGC light responses over large spatial scales (McIlwain, 1964; Olveczky et al., 2003). From such observations, a picture of retinal connectivity emerges in which all possible interactions between large numbers of cells must be probed before the function of the circuit can be understood (see Fig. 1A). Conversely, several lines of evidence indicate that synchronized firing can originate from pairwise, adjacent interactions: gap junction coupling between adjacent RGCs, electrical coupling of neighboring RGCs via an intermediate amacrine cell, or common synaptic inputs from bipolar or amacrine cells to neighboring RGCs (Mastrorarde, 1983b; Dacey and Brace, 1992; Jacoby et al., 1996; Brivanlou et al., 1998; Hu and Bloomfield, 2003; Hidaka et al., 2004; Schubert et al., 2005; Volgyi et al., 2005). Thus, interactions among multiple RGCs over large spatial scales may simply reflect the combined effect of pairwise, adjacent interactions (see Fig. 1D), which can be characterized using readily available experimental methods.

In this study, we test whether multi-neuron firing patterns are consistent with purely pairwise-adjacent connectivity or instead imply more complex circuitry. Until recently, two major challenges have precluded such an investigation. First, one must be able to record simultaneously from all cells in a circuit over a substantial area (Segev et al., 2004; Frechette et al., 2005). Second, one must have a framework within which to assess the significance of higher-order interactions among the neurons in a circuit (Amari, 2001; Schneidman et al., 2003b, 2006). We approach these issues with a combination of novel techniques. We apply new 512-electrode electrophysiological recording (Litke et al., 2004; Frechette et al., 2005) to the primate retina, which contains 1–2 dozen distinct RGC types that convey complete, parallel images of the visual scene to distinct targets in the brain and which closely resembles the human retina (Rodieck, 1998). In these recordings, distinct cell types such as the ON and OFF parasol cells are easily identified (see Fig. 2A) (Polyak, 1941; Watanabe and Rodieck, 1989). The regular mosaic arrangement of these cells in the retina (Peichl and Wässle, 1981; Wässle et al., 1983; DeVries and Baylor, 1997), combined with recordings that sample almost all of the cells of each type, make possible a direct test of pairwise and adjacent connectivity. We perform this test by examining the spatial patterns of electrical activity in different cells using the maximum entropy framework borrowed from statistical mechanics (Jaynes, 1957a,b), as described previously (Martignon et al., 2000; Amari, 2001; Schneidman et al., 2003b, 2006). The results indicate that patterns of electrical activity in groups of parasol cells may be understood almost entirely based on pairwise interactions, restricted to adjacent cells in the mosaic. This finding substantially simplifies our understanding of the retinal circuit.

Materials and Methods

Recordings. Preparation and recording methods were described previously (Chichilnisky and Kalmar, 2002; Litke et al., 2004; Frechette et al., 2005). Briefly, eyes were obtained from deeply and terminally anesthetized macaque monkeys (*Macaca mulatta*) used by other experimenters in accordance with institutional guidelines for the care and use of animals. Immediately after enucleation, the anterior portion of the eye and vitreous were removed in room light, and the eye cup was placed in a bicarbonate-buffered Ames’ solution (Sigma, St. Louis, MO) and stored in darkness at 32–34°C, pH 7.4, for ≥ 20 min before dissection. Under infrared illumination, pieces of peripheral retina 3–5 mm in diameter, isolated from the retinal pigment epithelium, were placed flat against a planar array of 512 extracellular microelectrodes, covering an area of $1800 \times 900 \mu\text{m}$. The present results were obtained from 30–60 m segments of recording. The preparation was perfused with Ames’ solution bubbled with 95% O_2 , 5% CO_2 and maintained at 32–34°C, pH 7.4.

Spike sorting. The voltage on each electrode was digitized at 20 kHz and stored for off-line analysis. Details of recording methods and spike sorting have been given previously (Litke et al., 2004). Briefly, spikes were identified using a threshold of three times the voltage SD. For each spike, the waveform of the spike and the simultaneous waveforms on six adjacent electrodes were extracted. Three to five waveform features were identified using principal component analysis. A mixture-of-Gaussians model was fit to the distribution of features using expectation maximization (Duda et al., 2001). The number of clusters and initial conditions for the model were determined automatically using an adapted watershed transformation (Castleman, 1996; Roerdink and Meijster, 2001). All clusters were visually inspected, and, when necessary, a mixture-of-Gaussians model was instead fit using manually selected initial conditions. Clusters with a large number of refractory period violations ($>10\%$ estimated contamination) or spike rates below 1 Hz were excluded from additional analysis.

Stimulation and receptive field analysis. An optically reduced stimulus

from a gamma-corrected cathode ray tube computer display refreshing at 120 Hz was focused on the photoreceptor outer segments. The low photopic intensity was controlled by neutral density filters in the light path. The mean photon absorption rate for the long (middle, short) wavelength-sensitive cones was approximately equal to the rate that would have been caused by a spatially uniform monochromatic light of wavelength 561 (530, 430) nm and intensity 9200 (8700, 7100) photons/ $\mu\text{m}^2/\text{s}$ incident on the photoreceptors. For the collection of parasol cells shown in Figure 2A, the mean firing rate during exposure to a steady, spatially uniform display at this light level was 10.7 ± 3.3 Hz for ON cells and 17.1 ± 3.5 Hz for OFF cells.

Spatiotemporal receptive fields were measured using a dynamic checkerboard (white noise) stimulus in which the intensity of each display phosphor was selected randomly and independently over space and time from a binary distribution. Root mean square stimulus contrast was 96%, and stimulus duration was 30 ms. The pixel size (60 μm) was selected to accurately capture the spatial structure of parasol cell receptive fields. For each RGC, the spike-triggered average stimulus was computed; this summarizes how the cell integrates visual inputs over space and time (Marmarelis and Naka, 1972; Chichilnisky, 2001). An elliptical two-dimensional Gaussian function was fit to the spatial profile; outlines in Figure 2A represent 1 SD boundary of these fits.

To analyze adjacent interactions, the mosaic structure of receptive fields was exploited. The modal distance between cells, D , was computed by examining a histogram obtained from all cell pairs by dividing the separation between cells in the pair by the mean 1 SD diameter of the cells and identifying the major mode near 0. For the ON and OFF parasol cells, this value was 2 (on average, 135 μm) and 2.4 (on average, 113 μm), respectively. Neighboring cells were identified as those for which the normalized distance was less than $D(1 + \sqrt{3})/2$, a value halfway between the nearest neighbor and next-nearest neighbor separation in a perfect triangular mosaic with elementary spacing D .

In testing the model restricted to pairwise-adjacent interactions (Fig. 1), note that the firing of two cells separated by an intermediate cell would be expected to reflect pairwise interactions via the intermediate cell. If the latter cell is excluded from analysis (that is, there is a “gap” in the group), the analysis will tend to reject the pairwise-adjacent model. This effect can also occur in a path of pairwise interactions through multiple intermediate cells; thus, it cannot be entirely eliminated without recording from the entire retina. To minimize this effect, groups of cells covering an approximately convex area of retina were selected for analysis. In the selection algorithm, a group of size N was generated by starting with a single cell and, on each iteration, adding a cell randomly selected from the collection of all cells having the maximum number of neighbors in the existing group, until the group had a total of N cells. To avoid problems in sparsely sampled regions of the mosaic, groups were rejected if they did not conform to two criteria. First, every cell in the group must have at least one neighbor in the group. Second, for every pair of adjacent cells in the group, there must be at least one other cell in the group adjacent to both cells of the pair.

Maximum entropy. Maximum entropy methods are used in statistical inference to identify an unknown distribution given several constraints that are insufficient to fully specify the answer. A parsimonious unique solution is to select the distribution with the greatest entropy consistent with the constraints (Jaynes, 1957a,b; Cover and Thomas, 1991). The entropy indicates the average number of bits required to transmit the identity of samples from a distribution (Cover and Thomas, 1991). Mathematically, maximizing the entropy is equivalent to selecting the maximum likelihood distribution consistent with the observed data (Berger et al., 1996).

Specifically, suppose the unknown distribution is $P(x)$, where x is a vector. Assume that N constraints on the distribution are known:

$$E[f_i(x)] = k_i, \quad \forall i \in \{1, \dots, N\}, \quad (1)$$

where $E[\cdot]$ is the expectation over the unknown distribution, and $f_i(\cdot)$ is

an arbitrary function of x . The maximum entropy solution is the distribution that maximizes entropy subject to the observed constraints:

$$\hat{P}(x) = \arg \max_{P(x)} \left[H[P] + \sum_{i=1}^N \lambda_i (E[f_i(x)] - k_i) \right], \quad (2)$$

where $H[P] = -\sum_x P(x) \log_2 P(x)$ is the entropy of a distribution $P(x)$, and λ_i are Lagrange multipliers. The second derivative matrix (or Hessian) of this equation is negative definite for all x . Thus, no local maxima exist and the unique global maximum can be found with any constrained gradient ascent optimization technique (Darroch and Ratcliff, 1972; Press et al., 1988; Berger et al., 1996; Martignon et al., 2000; Malouf, 2002). The functional form of the maximum entropy solution with pairwise constraints is an exponential distribution, which matches the Ising model in statistical mechanics (Hertz et al., 1991) and the Hopfield model of neural networks (Hopfield, 1982). The multipliers $\{\lambda_i\}$ measure the magnitude and sign of interactions between pairs of neurons; in the present work, the estimated multipliers were universally positive (data not shown).

For the present analysis, spike trains were binned at a resolution of 10 ms; bins with 2 or more spikes (<1%) were replaced with 1. The distribution $P(x)$ is over all binary words x expressed by M cells; this is a 2^M element vector (see Fig. 5). For example, for $M = 3$ cells, $P(x)$ is the probability distribution over the words $\{(0, 0, 0), (0, 0, 1), \dots, (1, 1, 1)\}$. The constraint is a linear function specifying which distribution values to sum to generate the observed marginal distributions, e.g., $P(1, 1) = P(1, 1, 1) + P(1, 1, 0)$. In matrix form, all constraints can be expressed as a $N \times 2^M$ sparse matrix, with values of 1 selecting which elements of the joint distribution sum to form the observed marginals. The exponential growth of the number of words with the number of cells analyzed necessitates efficient algorithms to estimate the maximum entropy distribution (Berger et al., 1996; Malouf, 2002) and effectively limits the number of cells that can be analyzed simultaneously. Any constrained gradient ascent method can be used to find $\hat{P}(x)$, although in practice, specialized algorithms exploiting Jensen’s inequality can be faster (Berger et al., 1996). The search terminated when the fractional change in entropy was smaller than 10^{-8} .

Likelihood analysis. Models were compared using likelihood analysis (see Fig. 6). The likelihood of a dataset under a given model is the probability of having observed that dataset under the assumption that the model is correct. In general, a more accurate model results in higher likelihood. In the present analysis, given a model probability distribution P over a set of firing patterns $\{f_i\}$ and given corresponding observed counts $\{c_i\}$ of each of these patterns, the log likelihood of the observed counts is as follows:

$$L = \log_2 \frac{(\sum_i c_i)!}{\prod_i c_i!} \prod_i P(f_i)^{c_i}, \quad (3)$$

where the leading term is a normalization to account for the fact that a given set of counts can arise from multiple different time series. Note that, in general, $P(f_i) < 1$; hence, longer data collection results in lower likelihood. Therefore, Figure 6 shows the average log likelihood, $\bar{L} = L/\sum_i c_i$. Using Stirling’s approximation for the leading term in Equation 3, it is easily shown that, for large datasets, $-\bar{L}$ is an estimate of the Kullback Leibler (KL) divergence D between the observed probability distribution P_{obs} , whose entries are $\{c_i/\sum_i c_i\}$ and the model probability distribution $\{P(f_i)\}$ (Cover and Thomas, 1991, their section 12.1). Using terminology from recent theoretical work, the index $(D_{\text{ind}} - D_{\text{pair}})/D_{\text{ind}}$ used to quantify the accuracy of the pairwise model is a cross-validated version of the fraction of the “multi-information” captured by the pairwise (or second-order) “connected information” (Schneidman et al., 2003b).

Estimating the KL divergence in finite data are nontrivial and subject to severe biases, especially in high-dimensional data (e.g., firing pattern distributions from large numbers of cells or over time). This estimation problem is intimately related to estimating the entropy of a distribution,

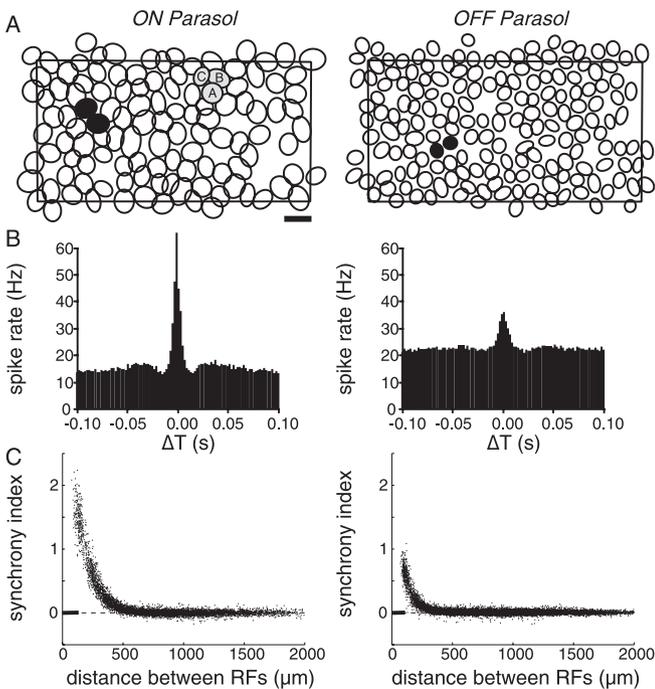


Figure 2. Pairwise synchrony. *A*, Receptive fields of 118 ON and 175 OFF parasol RGCs simultaneously recorded in one retina. Ellipses represent 1 SD of the Gaussian fit to the spatial profile of the receptive field. Rectangles represent outline of electrode array. Scale bar, 200 μm . *B*, Sample cross-correlation functions for pairs of ON and OFF parasol cells shown shaded in black in *A*. The firing rate of one cell is shown as a function of time relative to the time of a spike in the second cell. The synchrony index for these pairs is $S = 1.57$ and 0.59 , respectively. *C*, Synchrony index S as a function of distance between receptive fields of ON and OFF parasol cells. For comparison, the black bar near the origin represents the modal separation between cells in the mosaic (see Materials and Methods).

as well as applications in universal source coding, and is the subject of intense investigation (Nemenman et al., 2002; Paninski, 2003; Orlicsky et al., 2004). In simulation (data not shown), the bias of KL estimates was not significant for firing pattern distributions obtained from $M \leq 7$ cells, using standard techniques (Krichevsky and Trofimov, 1981). In the most extreme case, the mean occupancy of each bin in the firing pattern distribution was on the order of 100–1000, qualitatively indicating an estimate with minimal bias.

Results

Action potentials were recorded from RGCs in isolated macaque monkey retina perfused with physiological saline solution, using an array of 512 electrodes (Litke et al., 2004; Frchette et al., 2005). The receptive field of each RGC was identified using reverse correlation with a white-noise stimulus. Cells were segregated into distinct classes according to their receptive field characteristics; ON and OFF parasol cells were identified by their size and response kinetics (Chichilnisky and Kalmar, 2002). The receptive fields of all ON and OFF parasol cells in one recording are shown in Figure 2*A*.

These form an orderly mosaic that tiles visual space with minimal overlap, as expected from previous work in primates and other species (Wassle et al., 1983; Dacey, 1993; DeVries and Baylor, 1997). The completeness of the mosaics indicates that the recordings sampled most of the parasol cells in the $4 \times 8^\circ$ region of retina.

Responses of nearby RGCs are not statistically independent

To quantify functional connectivity in the retinal circuit, spontaneous activity of ON and OFF parasol cells was recorded in the

presence of steady, spatially uniform, photopic illumination. Interactions between pairs of cells were probed by examining the cross-correlation in spiking activity, that is, the probability that one cell fired as a function of time relative to the time of a spike in the second cell. If the two cells fired independently, the cross-correlation would be flat. Instead, nearby cells of the same type (ON or OFF) exhibited a strong tendency to fire nearly synchronously, as shown by the pronounced peak at the origin with a width of ~ 10 ms (Fig. 2*B*). This synchronized firing in the absence of time-varying stimulation (also known as “noise correlations”) resembles observations in other species and indicates functional connectivity attributable to common inputs and/or reciprocal connections in the retina (Arnett, 1978; Johnsen and Levine, 1983; Mastronarde, 1983b; Meister et al., 1995; DeVries and Baylor, 1997; Brivanlou et al., 1998; Hidaka et al., 2004). Anticorrelation between ON and OFF parasol cells was also observed (data not shown) (Mastronarde, 1983a), but subsequent analysis will be restricted to cells of the same type.

Pairwise synchrony was restricted to nearby cells. To quantify this tendency, spike trains from pairs of cells were binned at a resolution of 10 ms, yielding binary spike counts A and B as a function of time. The probability of synchronous firing, $P(A = 1, B = 1)$, was measured, where 1 denotes the occurrence of a spike. This probability was expressed relative to the probability expected from statistically independent firing using a “synchrony index:”

$$S = \log_2 \frac{P(A = 1, B = 1)}{P(A = 1)P(B = 1)} \quad (4)$$

The index is shown as a function of separation between pairs of cells in Figure 2*C*. Cells separated by less than several hundred micrometers, corresponding to approximately three times the average cell–cell separation in the mosaic, fired synchronously at rates well above chance ($S = 0$). Cells that were more widely separated did not. In summary, pairs of parasol RGCs in primate retina are functionally connected, and this connectivity is universal among cells of like type and spatially localized (Mastronarde, 1983b).

The above observations provide significant constraints on functional connectivity between RGCs. However, the connectivity of the circuit cannot be understood with these measurements alone. First, the fact that pairwise synchrony extends over a distance of several cells in the mosaic (Fig. 2*C*) does not require that direct connections reach this far. Instead, synchrony might be mediated by pairwise interactions, for example, gap junction electrical coupling, propagated through several intermediate cells (Dacey and Brace, 1992; Jacoby et al., 1996; Hu and Bloomfield, 2003; Hidaka et al., 2004; Schubert et al., 2005; Volgyi et al., 2005). Second, the observed pairwise synchrony provides no information about higher-order interactions. To illustrate the distinction between pairwise and higher-order synchrony, consider two hypothetical mechanisms: (1) common input from a spiking amacrine cell to more than two RGCs (Meister et al., 1995; Schnitzer and Meister, 2003), and (2) direct gap junction electrical coupling between pairs of RGCs (note that these candidate mechanisms are not exhaustive but merely serve as examples). These mechanisms can produce very similar cross-correlations in pairs of cells but very different concerted firing patterns in multiple cells. Specifically, in case 1, a single spike in an amacrine cell could induce a spike in most or all of its RGC targets simultaneously. However, in case 2, the probability of the entire collection of RGCs firing synchronously is relatively low. The distinc-

tion is fundamental: the mechanism of case 1 has been suggested as a way to convey distinctive, multiplexed visual messages to the brain (Meister et al., 1995; Schnitzer and Meister, 2003). This simple example illustrates that pairwise synchrony provides only a limited picture of circuit connectivity. Thus, probing the full extent of neural interactions requires measuring patterns of firing in many cells simultaneously.

Triplet synchrony is explained by pairwise interactions

Consider three adjacent ON parasol cells labeled A, B, and C in Figure 2A. Each cell pair exhibits strong synchronized firing (Fig. 3A). To characterize the joint activity of the three cells requires extending the notion of cross-correlation to three dimensions. The left plot in Figure 3B shows the probability of spikes in all three cells as a function of time shifts ΔT between cells A and B and between cells A and C. The three extended ridges represent the preponderance of synchronous firing in each of the cell pairs, regardless of the activity of the third cell, similar to Figure 3A (for an alternative representation, see Perkel et al., 1975). However, the prominent peak at the origin, with a width of ~ 10 ms, indicates that all three cells often fired synchronously. The key question is whether this multi-cell firing occurs at the rate expected from the known pairwise correlation (e.g., mechanistic hypothesis 2 above) or, alternatively, whether more complex patterns of interaction are required to explain triplet firing (e.g., mechanistic hypothesis 1 above).

To distinguish these possibilities requires a null hypothesis: the probability of all three cells firing in the same time bin ($\Delta t = 10$ ms), given only the pairwise interactions. This is analogous to the two-cell case, in which the null hypothesis was statistical independence and a synchrony index much greater than 0 revealed a departure from the null hypothesis. However, in the case of three or more cells, the null hypothesis is not as simple, because the statistical independence hypothesis,

$$P_{\text{ind}}(A,B,C) = P(A)P(B)P(C), \quad (5)$$

does not account for known pairwise correlations. Ideally, one would compare the observed data with a null hypothesis that assumes nothing beyond pairwise correlations.

An approach borrowed from statistical mechanics, known as “maximum entropy,” solves the problem uniquely (Jaynes, 1957a,b; Berger et al., 1996; Martignon et al., 2000; Amari, 2001; Scheidman, 2003b) (see Materials and Methods). Intuitively, the approach is to pose a null hypothesis that assumes no prevalence of structure in firing patterns above and beyond the observed pairwise interactions. Specifically, denote the joint distribution of responses of the three cells by $P(A,B,C)$, the pairwise joint distributions by $P(A,B)$, $P(B,C)$, and $P(A,C)$, and the single-cell firing rates by $P(A)$, $P(B)$, and $P(C)$. Our null model, $P_{\text{null}}(A,B,C)$, must satisfy the single-cell and pairwise constraints. This problem is underconstrained because many different models for the joint activity are consistent with these marginal constraints. The

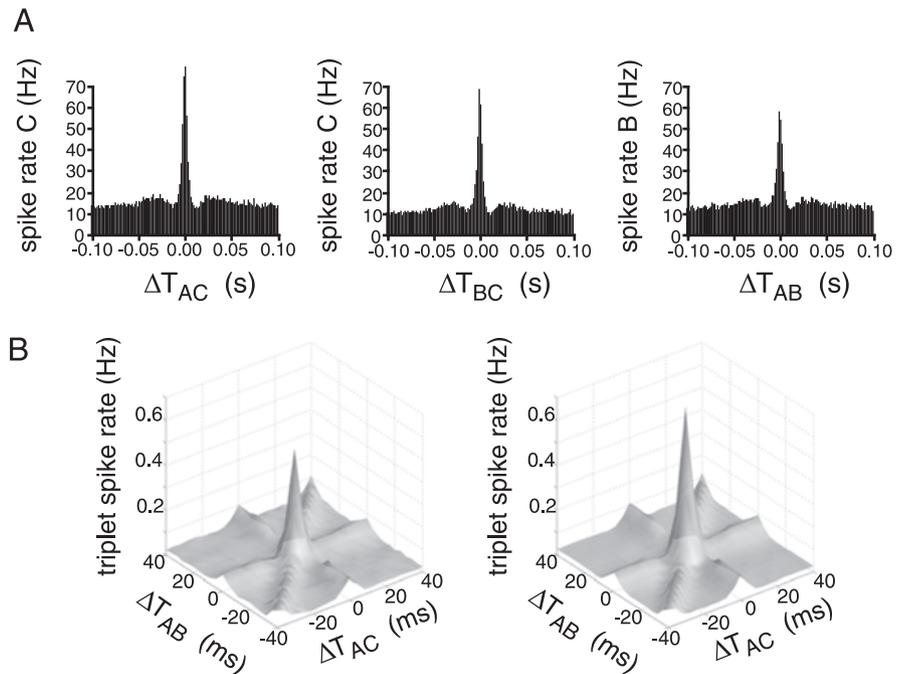


Figure 3. Triplet synchrony. **A**, Cross-correlation functions for each pair from a group of three adjacent ON parasol cells (labeled A, B, C in Fig. 2A). Each cell pair exhibits strong synchrony ($S = 1.90, 2.05$, and 1.64 from left to right). **B**, Three-dimensional cross-correlation for triplet synchrony (left) and maximum entropy pairwise prediction (right). Note that maximum entropy predictions computed for different time offsets are not statistically independent.

solution is to compute a joint distribution, $P_{\text{pair}}(A,B,C)$, that has the maximum possible entropy (i.e., least structure) given that it satisfies the marginal constraints. $P_{\text{pair}}(A,B,C)$ is a complete null hypothesis for the observed firing patterns. Note that, in the two-cell case, maximum entropy is equivalent to statistical independence: $P_{\text{null}}(A,B) = P_{\text{ind}}(A,B)$. The right panel in Figure 3B shows the maximum entropy pairwise prediction for all three cells firing at various time offsets. This model accounts for the major features of the three-dimensional cross-correlation, qualitatively consistent with the idea that pairwise interactions alone can explain the temporal structure of synchronized firing in these three cells.

To test the maximum entropy pairwise prediction quantitatively, analysis was restricted to patterns of firing in the same time bin for all three cells (e.g., central peaks in Fig. 3B). Although this analysis does not account for the possibility of complex temporal interactions, it captures the spatial patterns of synchronous activity that dominate both the pairwise and triplet cross-correlograms. A “pattern index” was computed that generalizes the synchrony index (Eq. 4) to all spike patterns in three cells:

$$Q = \log_2 \frac{P_{\text{obs}}(A,B,C)}{P_{\text{null}}(A,B,C)}. \quad (6)$$

In this expression, $P_{\text{obs}}(A,B,C)$ denotes the observed probability of a particular three-cell firing pattern. P_{null} represents a null hypothesis for the prevalence of different firing patterns. The index Q indicates how frequently a particular firing pattern was observed, relative to the prediction of the null model P_{null} . Note that, in the reduced case of two cells firing at the same time ($A = B = 1$) and a null model of statistical independence ($P_{\text{null}} = P_{\text{ind}}$), Q is equivalent to the index S in Equation 4.

First, the degree of triplet synchronized firing was established

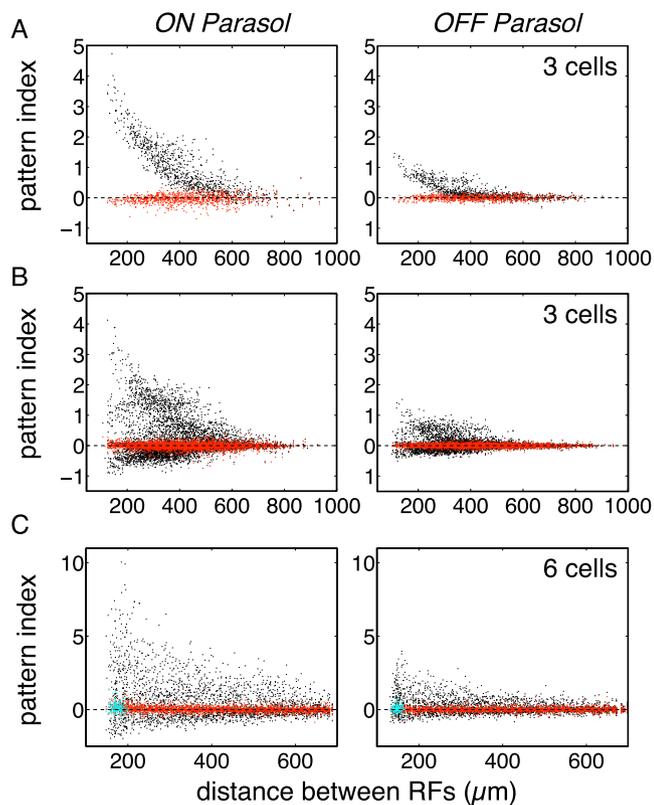


Figure 4. Pattern index for multi-neuron firing. **A**, Pattern index Q (Eq. 6) for firing pattern (1, 1, 1) in groups of three cells as a function of the geometric mean of the distances between each pair in the group. Black symbols, Pattern index using statistical independence as the null model. Red symbols, Pattern index using pairwise constrained maximum entropy as the null model. **B**, Pattern index for all eight firing patterns in groups of three cells. Symbol colors as in **A**. **C**, Pattern index for groups of six cells. Red and black symbols as in **A**. Blue symbols, Pattern index obtained using the pairwise-adjacent maximum entropy model. For the latter, analysis was restricted to local groups of cells conforming to selection criteria described in Materials and Methods and containing at least one non-adjacent cell pair.

by examining deviations from statistical independence. For this analysis, $P_{\text{null}} = P_{\text{ind}}$. The black symbols in the left panel of Figure 4A show the value of Q computed for the firing pattern (1, 1, 1) in many cell triplets randomly sampled from the mosaic of ON parasol RGCs shown in Figure 2A. The index is plotted as a function of the geometric mean of the distances between each pair of cells in the triplet. As with the pairwise synchrony index (Fig. 2C), Q is large for nearby triplets and approaches 0 for widely spaced triplets. Specifically, triplets of parasol cells fired synchronously up to $2^Q \approx 10$ times more frequently than expected from statistical independence. Note that triplets with a large mean distance can sometimes exhibit departures from independence if two of the cells in the triplet are near one another. Similar results are shown for OFF parasol cells in the right panel.

Second, to assess whether pairwise interactions alone account for triplet synchrony, the predictions of the maximum entropy model using pairwise constraints were examined. In this case, $P_{\text{null}} = P_{\text{pair}}$. The values of Q for this model (red symbols) cluster near the horizontal line at 0, that is, the probability of triplet synchrony is approximately equal to the probability expected from pairwise interactions. Thus, the excess of triplet firing above what is expected from statistical independence is accounted for by pairwise interactions.

All triplet firing patterns are explained by pairwise interactions

To test the maximum entropy pairwise prediction more fully, analysis was performed on all simultaneous firing patterns for triplets of recorded cells. For example, all cells firing synchronously is represented by (1, 1, 1), and no firing is represented by (0, 0, 0). The distribution of all eight possible firing patterns observed in a set of three ON parasol cells is shown in Figure 5A. For comparison, Figure 5C shows the firing pattern distribution expected from statistical independence (Eq. 5), which deviates substantially from the data. Figure 5B shows the maximum entropy pairwise prediction, which captures most of the structure in the data.

To quantify these observations for each firing pattern separately, values of Q were examined for all eight possible triplet firing patterns. In Figure 4B, distinct values of Q for each firing pattern in a triplet of cells are shown with different points. The band of points near the top reproduces the data of Figure 4A, that is, the firing pattern (1, 1, 1). Other patterns occur with greater or lesser frequency than predicted from statistical independence because strongly coupled cells tend to fire together or not at all (Fig. 5A, C). The departure from independence in all triplet firing patterns, evidenced by the large vertical spread of black points, is almost entirely eliminated by taking into account pairwise interactions (red symbols).

To provide a statistical summary of the performance of the pairwise model across all firing patterns, the average likelihood of the observed data, P_{obs} , under different models was computed. The likelihood indicates the probability of having observed the data given a particular model. In general, more (less) accurate models exhibit likelihood values approaching 1 (0). To express the likelihood in units that are invariant with respect to the length of the data sample, the average (geometric mean) likelihood per time bin was computed (see Materials and Methods). The average likelihood of the data in Figure 5A under the independent model (Fig. 5C) was 0.94959, whereas the likelihood under the pairwise model (Fig. 5B) was 0.99944. Although this difference seems modest, the multiplicative accumulation of probability over sequential time bins means that, from 60 s of data collection, one may conclude that the probability of the data having arisen from the pairwise model is $>10^{133}$ times larger than the probability of the data having arisen from the independent model. In other words, the enormous difference between the predictive power of the pairwise and independent models is easily distinguished even with modest data sets.

As a benchmark for the likelihood, an empirical model P_{emp} was obtained from the recorded frequencies of all eight firing patterns in a separate segment of recording from the same retina (the same recording that was used to fit P_{ind} and P_{pair} above). P_{emp} represents an ideal model for the observed data P_{obs} given the intrinsic reproducibility of the measurement. For the data in Figure 5A, the average empirical likelihood was 0.99955. Compared with the ratio of likelihoods above, the ratio of empirical to pairwise model likelihood was smaller than 2 with 60 s of data. Note that the empirical likelihoods of the data are not 1, reflecting a combination of finite counting statistics and possible (small) nonstationarity in the neurophysiological recording. Also, note that the empirical likelihood is an upper bound for the performance of any model; as expected, it exceeds the likelihood associated with both the pairwise and independent models. The pairwise model likelihood is very close to this upper bound, whereas the independent model likelihood is far from it.

The average likelihood was examined for all triplets of cells

recorded. The dark blue symbols in Figure 6A show the likelihood of the observed triplet firing patterns assuming statistical independence, P_{ind} , compared with the likelihood obtained with the empirical model, P_{emp} . Again, the likelihood under the independent model is much lower, confirming that the firing of RGCs departs substantially from statistical independence. In contrast, Figure 6B shows that the pairwise model P_{pair} produces likelihood values nearly identical to those produced by the empirical model P_{emp} . Thus, pairwise interactions explain the frequency of triplet firing patterns nearly as accurately as a repeated measurement.

Multi-cell firing patterns are explained by pairwise interactions

The maximum entropy framework is easily extended to test for interactions in larger groups of cells. For example, for five-cell firing patterns, the observed firing pattern distribution, $P_{obs}(A, B, C, D, E)$, is compared with a maximum entropy null model, $P_{pair}(A, B, C, D, E)$, and a statistically independent model, $P_{ind}(A, B, C, D, E)$. Figure 5D–F shows these three firing pattern distributions for collections of five cells. As in the three-cell case, the pairwise model prediction substantially captures the structure of the observed firing pattern distribution, whereas the statistically independent prediction does not.

As above, the pattern index was examined to quantify the quality of the predictions. Figure 4C shows values of Q for the independent and pairwise models, for groups of six cells randomly sampled from the ON and OFF parasol mosaics of Figure 2A. Black symbols reveal large departures from statistical independence: values of Q as high as ~ 10 indicate firing patterns occurring ~ 1000 more frequently than predicted by chance. Red symbols reveal that these departures are almost entirely accounted for by pairwise interactions. This finding is supported by comparing the likelihood of the observed firing pattern distribution P_{obs} under the models P_{ind} and P_{pair} with the likelihood obtained from the empirical model P_{emp} for groups of four, five, six, and seven cells. [For this analysis and what follows, results from larger groups are not reported because of potential biases attributable to high dimensionality of the data (see Materials and Methods).] In all cases, the large systematic failures of statistical independence (Fig. 6A) are explained by the pairwise model (Fig. 6B).

In summary, pairwise interactions explain almost all of the departures from statistical independence in parasol cell signals, with a precision comparable with the reproducibility of the measurements. This implies that the structure of multi-neuron firing patterns may be understood with high accuracy based on pairwise connectivity, without postulating more complex interactions. A possible mechanistic interpretation is that frequent multi-cell firing in the retina does not imply widely diverging common inputs (Schnitzer and Meister, 2003) but instead can arise from reciprocal connections between, or common inputs to, pairs of cells (see Discussion).

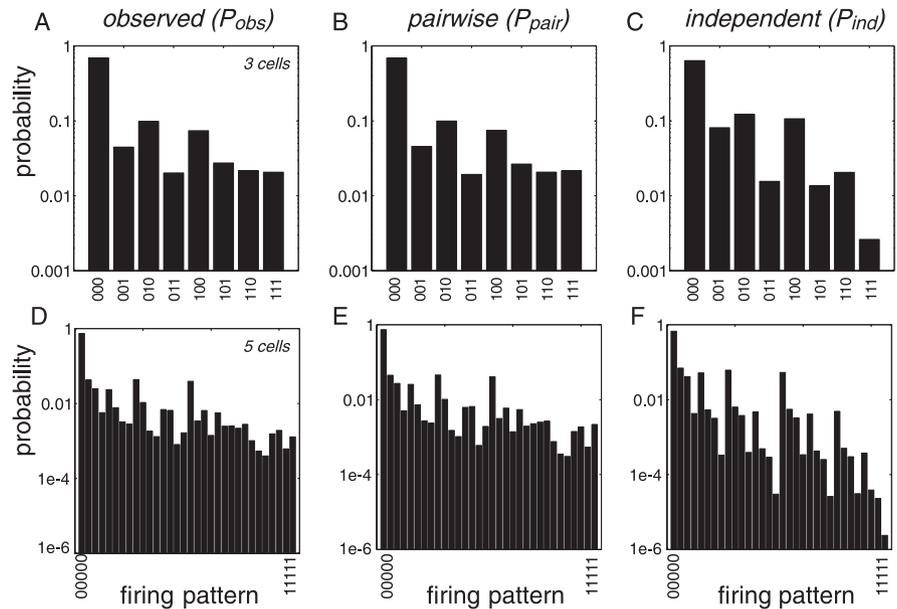


Figure 5. Sample firing pattern distributions. **A–C** show firing pattern distributions for three ON parasol cells. Each distribution is shown as a histogram, with the firing pattern indicated by binary digits on the abscissa representing a spike or no spike in each cell, and the probability indicated on the ordinate (note logarithmic scale). Firing pattern distributions were obtained from raw data (**A**), the pairwise model (**B**), and the statistically independent model (**C**) (the latter obtained by multiplying the marginal probability distributions for each cell; see Eq. 5). The likelihood of the data obtained from the pairwise model was 0.99944, from the independent model was 0.94959, and from the empirical model was 0.99955. Distributions in **B** and **C** were fitted to interleaved recordings from the same cells, distinct from the data in **A** (data not shown). **D–F**, Same as **A–C** but for five ON parasol cells. The likelihood of the data obtained from the pairwise model was 0.99728, from the independent model was 0.92107, and from the empirical model was 0.99836.

Multi-cell firing patterns are explained by pairwise, adjacent interactions

The next step in simplifying our understanding of the circuit is to test whether pairwise interactions are restricted to immediately adjacent cells or instead must span larger distances. The mosaic arrangement of receptive fields, which corresponds to the physical arrangement of RGCs in the retina (Wassle et al., 1983; DeVries and Baylor, 1997), provides a uniquely useful measure of adjacency (Fig. 2A), and the maximum entropy approach provides a straightforward test of the hypothesis.

Specifically, consider a more restrictive null model for the distribution of firing patterns, P_{adj} . This is the unique maximum entropy distribution subject to the single-cell constraints $\{P(A), P(B), \dots\}$ as before and pairwise constraints $\{P(A, B), P(B, C), \dots\}$ obtained only from immediately adjacent cells in the mosaic (see Materials and Methods). The pattern index obtained from many cell groups with this model are shown with blue symbols in Figure 4C. To reduce the effect of adjacent interactions present in the retina but missing in the group of cells analyzed, only groups that uniformly cover an approximately convex area of the retina were considered (see Materials and Methods). As a result, the points cover only a small range of distances. Trivial cases in which all pairs are adjacent were excluded. As with the more general pairwise model, the pairwise-adjacent model (turquoise symbols) accounts for almost all of the departures from statistical independence (black symbols). Figure 6C shows that the likelihood of the data P_{obs} under the pairwise-adjacent model P_{adj} is very similar to the likelihood under the empirical model from a repeated measurement P_{emp} . Comparison with Figure 6B shows that the pairwise and pairwise-adjacent model likelihoods are nearly indistinguishable. In summary, pairwise interactions be-

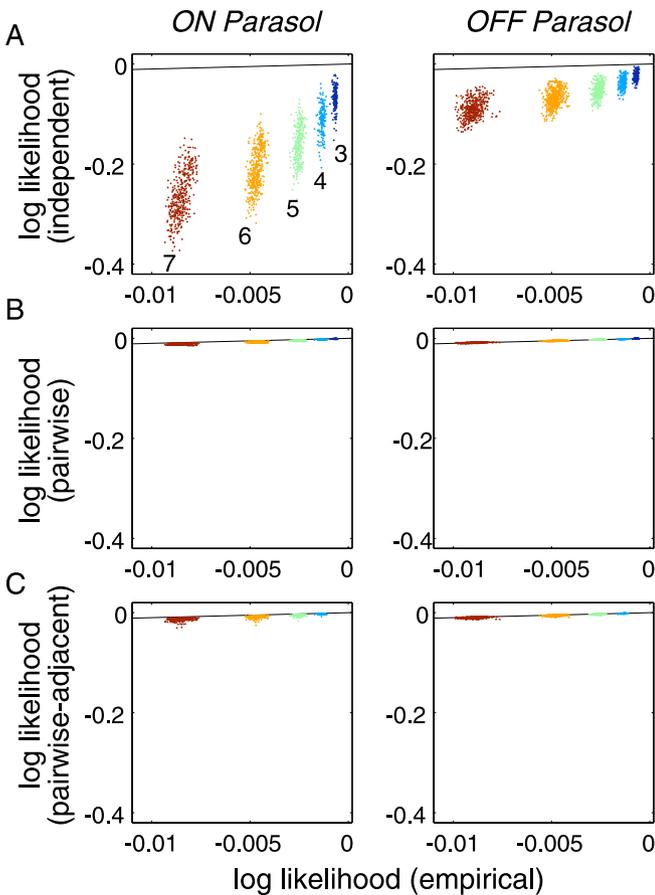


Figure 6. Likelihood test of pairwise and adjacent models under constant, spatially uniform illumination. **A**, Likelihood of observed data under an assumption of statistical independence as a function of likelihood obtained from an empirical model based on a repeated measurement, for groups of three, four, five, six, and seven cells. Diagonal gray line near top indicates equality (note different scales on abscissa and ordinate); large departures from this line indicate substantial failures of statistical independence. **B**, Likelihood of observed data in the pairwise maximum entropy model as a function of empirical likelihood. Symbol colors and equality line same as in **A**. **C**, Likelihood of observed data in the pairwise–adjacent maximum entropy model as a function of empirical likelihood. Symbol colors and equality line same as in **A**. All likelihood analysis was restricted to local groups of cells conforming to selection criteria described in Materials and Methods. For **C**, analysis was further restricted to groups that included at least one nonadjacent cell pair.

tween adjacent RGCs in the mosaic are sufficient to account for almost all multi-neuron firing patterns.

These findings imply that firing patterns in parasol cells can be understood with high accuracy using the simplest possible connectivity illustrated in Figure 1*D*, a dramatic reduction in complexity. A possible mechanistic interpretation of this finding is that pairwise synchrony between non-adjacent cells in the mosaic (Fig. 2*B*) do not necessarily imply long-distance contacts but instead may be explained by propagation via pairwise connections between adjacent cells (e.g., gap junctions). To test the propagation hypothesis, the synchrony index observed in pairs of cells was compared with the index predicted from the pairwise–adjacent model fitted to cell groups of size $n = 7$ (Fig. 8). Over the entire range of distances between cells, and across the range of synchrony index values observed, the observed and predicted synchrony indices were mostly similar, for both adjacent cell pairs and nonadjacent cell pairs. This is significant because, in nonadjacent cells, the pairwise–adjacent model only predicts synchrony as a consequence of synchrony with intermediate cells. Note, however, that systematic discrepancies are present, partic-

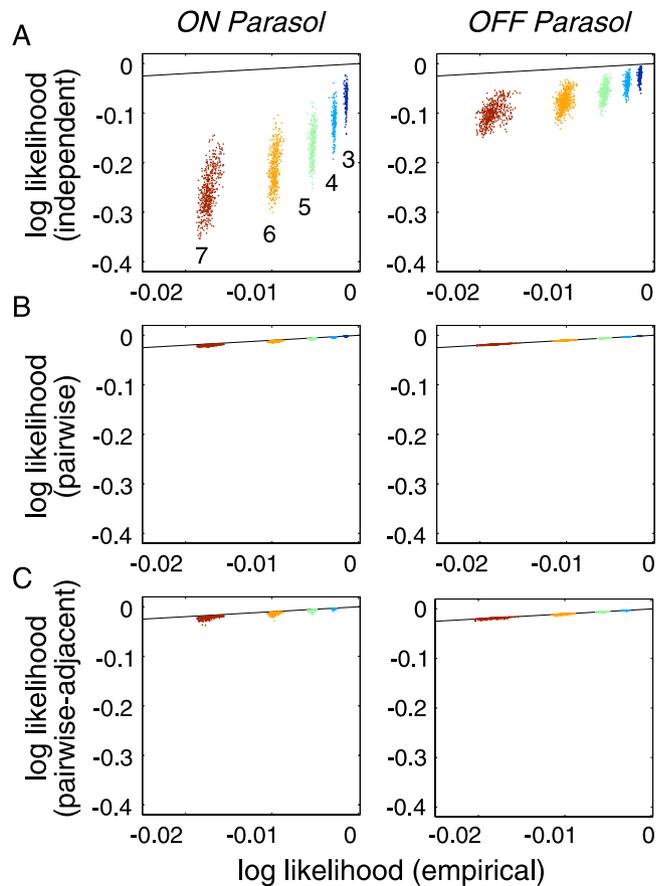


Figure 7. Likelihood test of pairwise and adjacent models with a fine-grained random visual stimulus. Panels and symbols as in Figure 6.

ularly for nonadjacent OFF cells, suggestive of subtle departures from the pairwise–adjacent model.

Measuring the accuracy of pairwise and pairwise–adjacent models

To quantify how accurately the pairwise and pairwise–adjacent models explain interactions between RGCs, the average log likelihood value \bar{L} shown in Figure 6 provides a natural measure. First, larger values of \bar{L} indicate that the observed data are more consistent with the model. Second, it is easily shown (see Materials and Methods) that the value $-\bar{L}$ is an estimate of the Kullback–Leibler divergence, D , between the observed firing pattern distribution, P_{obs} , and the model distribution. The divergence is an information-theoretic quantity that measures the inefficiency of storing the observed firing patterns using a compression scheme optimized for the model probability distribution (Cover and Thomas, 1991). Thus, larger divergence values (smaller values of \bar{L}) correspond to a less accurate model.

The divergence of the independent model, D_{ind} , quantifies the departures from statistical independence in the firing of different RGCs. A natural measure of the success of the models is the degree to which they capture these departures from independence. The index $(D_{\text{ind}} - D_{\text{pair}})/D_{\text{ind}}$ expresses the fraction of the departures from independence accounted for by the pairwise model. On average, this value was $\sim 99\%$ (Table 1). Similarly, the index $(D_{\text{ind}} - D_{\text{adj}})/D_{\text{ind}}$ expresses the fraction of departures from independence accounted for by the pairwise–adjacent model. On average, this value was $\sim 98\%$. As a benchmark, the index for the empirical model, $(D_{\text{ind}} - D_{\text{emp}})/D_{\text{ind}}$, was $\sim 99\%$.

The latter quantity represents the highest value that can be expected given the reproducibility of the data. The departure from independence accounted for by each model was $\sim 99\%$ of this benchmark value. In summary, both models almost entirely account for the departures from independence observed in RGC firing.

To test whether these results depend strongly on the timescale of the analysis (10 ms time bins), analysis was repeated with bin sizes twofold larger and smaller. The results in Table 1 indicate that the predictive power of pairwise and pairwise-adjacent models is essentially constant across this range of time scales.

Multi-cell firing patterns in the presence of visual stimulation are explained by pairwise, adjacent interactions

The data presented so far were obtained with steady spatially uniform illumination of the retina and thus reflect the circuitry that mediates spontaneous synchronized firing in RGCs. However, it is possible that some contributions to synchronized firing, such as diverging inputs from amacrine cells, arise only in the presence of visual stimuli that vary over space and time. To test this possibility, the maximum entropy analysis was applied to data collected in the presence of the white-noise stimulus used for the receptive field measurement (Fig. 7). This stimulus provides a wide range of spatial and temporal variations in an experiment of reasonable duration. Note that stimuli with a large spatial scale would be expected to introduce higher-order correlations by simultaneously activating multiple RGCs. This would confound the analysis because synchronized firing could be produced by the stimulus, the retinal circuitry, or both (Schneidman et al., 2003a). This problem was avoided by using a stimulus with independently modulating pixels that were small relative to the parasol cell receptive field. In these conditions, the pairwise model captured $\sim 98\%$ of the departures from independence, the pairwise-adjacent model captured $\sim 98\%$, and the empirical model benchmark captured $\sim 99\%$ (Table 1). Again, the models accounted for $\sim 99\%$ of the departures from independence that were reproduced by the empirical benchmark. Thus, even in the presence of a dynamic, spatially varying stimulus, pairwise and adjacent models almost entirely account for the departures from independence observed in RGC firing.

Sensitivity of maximum entropy analysis for detecting non-pairwise, non-adjacent circuitry

Pairwise and adjacent connectivity clearly can explain most of the observed interactions between parasol RGCs. However, without direct experimental manipulation, it is impossible to draw firm

Table 1. Accuracy of pairwise and pairwise-adjacent models

Condition	Cell type	Δt	Pairwise	Pairwise-adjacent	Empirical
Constant illumination	ON parasol	10 ms	$98.6 \pm 0.5\%$	$98.3 \pm 1.0\%$	$99.8 \pm 0.1\%$
		5 ms	$97.8 \pm 0.5\%$	$97.5 \pm 1.0\%$	$99.8 \pm 0.1\%$
		20 ms	$99.2 \pm 0.4\%$	$98.9 \pm 0.9\%$	$99.7 \pm 0.1\%$
	OFF parasol	10 ms	$98.9 \pm 0.3\%$	$98.2 \pm 1.3\%$	$99.4 \pm 0.4\%$
		5 ms	$98.5 \pm 0.4\%$	$97.4 \pm 1.5\%$	$99.0 \pm 0.5\%$
		20 ms	$99.2 \pm 0.3\%$	$98.6 \pm 1.3\%$	$99.4 \pm 0.3\%$
White-noise stimulus	ON parasol	10 ms	$98.6 \pm 0.4\%$	$98.4 \pm 0.9\%$	$99.6 \pm 0.2\%$
		5 ms	$97.8 \pm 0.5\%$	$97.5 \pm 1.0\%$	$99.6 \pm 0.2\%$
		20 ms	$99.1 \pm 0.3\%$	$98.9 \pm 0.9\%$	$99.5 \pm 0.2\%$
	OFF parasol	10 ms	$98.6 \pm 0.5\%$	$98.3 \pm 0.9\%$	$98.8 \pm 0.7\%$
		5 ms	$98.2 \pm 0.8\%$	$98.0 \pm 1.3\%$	$98.4 \pm 1.0\%$
		20 ms	$98.6 \pm 0.7\%$	$98.3 \pm 1.0\%$	$98.6 \pm 0.9\%$

Each numerical entry indicates the percentage of departures from statistical independence captured by a specific model, for a given stimulus condition, cell type tested, and time bin size (Δt). Results cited in text were obtained with $\Delta t = 10$ ms. The range indicates the mean ± 1 SD across several hundred cell groups tested, each consisting of three to seven cells (in the case of the pairwise model) or four to seven cells (in the case of the pairwise-adjacent model). For each model, the quantities shown are as follows: pairwise model, $(D_{\text{ind}} - D_{\text{pair}})/D_{\text{ind}}$; pairwise-adjacent model, $(D_{\text{ind}} - D_{\text{adj}})/D_{\text{ind}}$; empirical model, $(D_{\text{ind}} - D_{\text{emp}})/D_{\text{ind}}$.

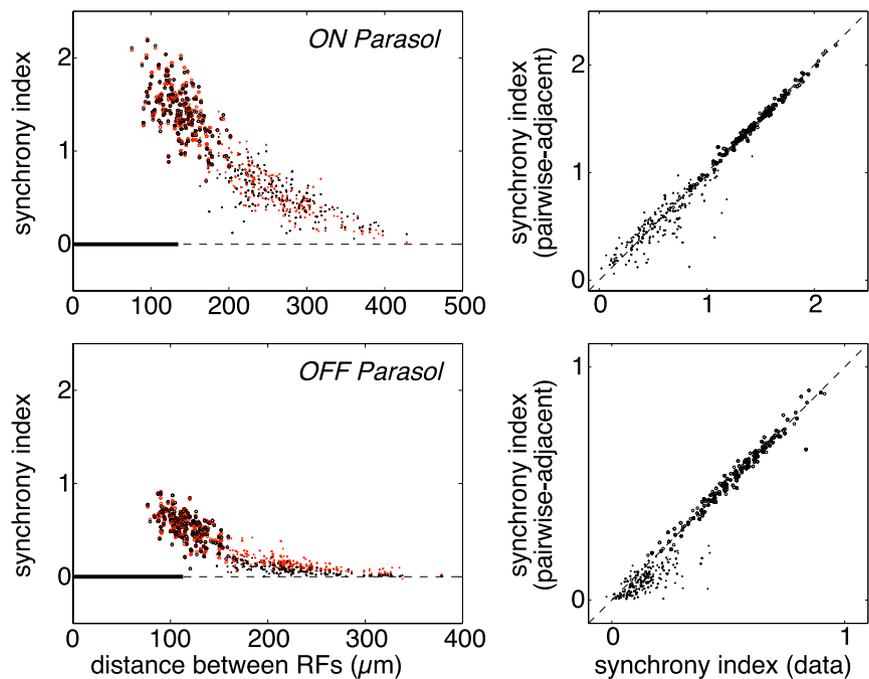


Figure 8. Predicted and observed pairwise synchrony index. Left panels show the synchrony index (Eq. 4) in pairs of parasol cells as a function of distance between their receptive fields (see Fig. 2). Black points indicate the observed synchrony index, and red points indicate the predictions obtained from the maximum entropy pairwise-adjacent model fitted to groups of $n = 7$ cells. The black bar near the origin represents the modal separation between cells in the mosaic. Right panels show the comparison between data and model predictions for each cell pair tested. In all panels, large open symbols represent cell pairs that are adjacent in the mosaic; small symbols represent cell pairs that are not.

conclusions about the structure of retinal circuits that underlies this statistical observation. A first step, however, is to ask, how sensitive is the maximum entropy analysis to departures from pairwise and adjacent connectivity? This question was approached by testing the sensitivity of the analysis to artificial perturbations of the data, in two ways.

First, the analysis was applied to artificial data obtained from a purely pairwise (or pairwise-adjacent) model, with varying amounts of common input added to simulate departures from the model. Specifically, a maximum entropy pairwise distribution, P_{pair} , was fitted to the observed data from a group of $n = 7$ RGCs. A hypothetical common input to all n cells was then simulated, occurring randomly with probability $0 < r < 1$ in each

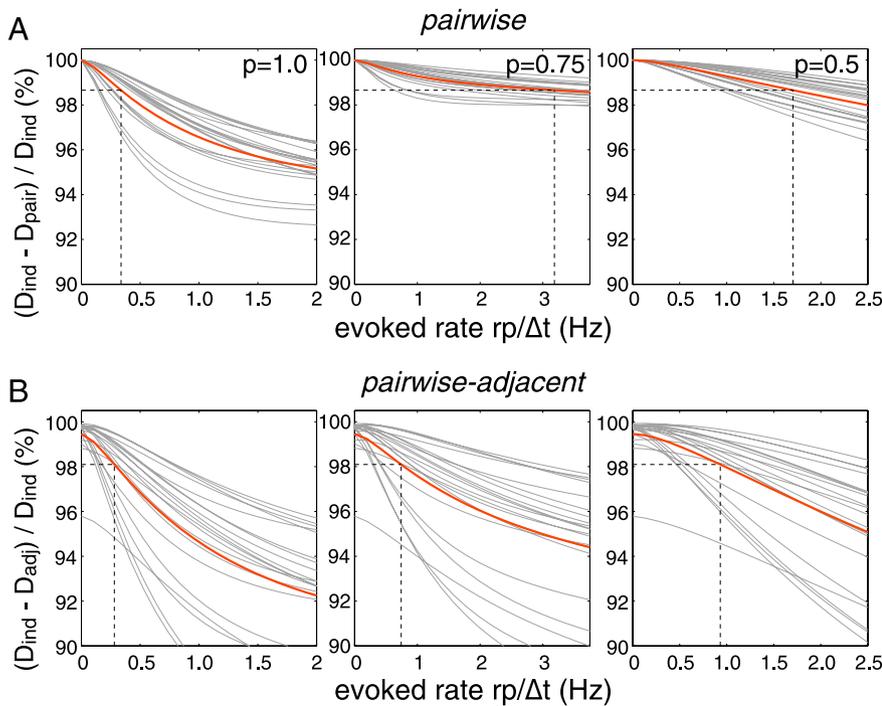


Figure 9. Sensitivity of maximum entropy analysis. Each panel shows the sensitivity of the maximum entropy analysis procedure for detecting a hypothetical non-pairwise, non-adjacent common input. The hypothetical input occurs at a rate r and causes a spike in each of $n = 7$ RGCs with a probability p . Results are shown for common input over a range of values of r and p added to simulations obtained with either the pairwise (**A**) or pairwise-adjacent (**B**) model fitted to data. In each case, the abscissa indicates the average evoked rate in the simulated RGCs, which is determined by r and p . The ordinate indicates the fraction of the departures from statistical independence accounted for by the pairwise or pairwise-adjacent model. Each gray trace shows the results obtained from a single group of ON parasol cells; red traces indicate the average across all 20 groups. In each panel, the rate of common input r required to reproduce the average value observed in the original data was converted to the equivalent evoked rate and is indicated by a dashed line.

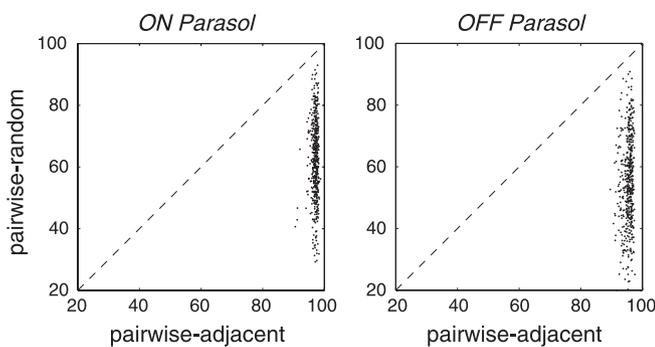


Figure 10. Sensitivity of maximum entropy analysis for detecting non-adjacent interactions. Each panel shows the fraction of departures from independence accounted for by the pairwise analysis restricted to a random subset of cell pairs, equal in number to the number of adjacent cell pairs, as a function of the fraction accounted for by the original pairwise-adjacent analysis. Each point represents the results for a single groups of $n = 7$ ON or OFF parasol cells. Dashed lines indicate equality.

time bin, and generating a spike in each RGC with efficacy expressed as a probability p . If (x_1, \dots, x_n) is the binary firing pattern for n cells, then define $m = \sum_i x_i$ to be the total number of spikes in the firing pattern. The common input alone would produce a particular RGC firing pattern containing m spikes with probability:

$$P_{\text{common}} = \begin{cases} rp^m(1-p)^{n-m} & m > 0 \\ (1-r) + r(1-p)^n & m = 0 \end{cases} \quad (7)$$

P_{common} is the probability distribution of firing patterns created by the nonpairwise common input. The effect of this common input was simulated by calculating the distribution of firing patterns expected from the logical or combination of spikes from the firing patterns produced by the common input distribution, P_{common} , and the pairwise distribution, P_{pair} . This resulting simulated firing pattern distribution contains structure with systematic departures from pairwise interactions.

To assess the sensitivity of the pairwise model to common input, the degradation in performance of the model was measured as a function of the putative input rate r in the simulation. Specifically, a new pairwise model was fitted to the simulated data, and the fraction of departures from statistical independence accounted for was computed. In the case of $r = 0$ (no common input), this value was 100%, because the pairwise model computed from the simulated distribution exactly reproduces the original P_{pair} . As Figure 9 demonstrates, the pairwise model is indeed sensitive to common input: this value systematically drops from 100% as the input rate r increases.

To assess the possible degree of common input in the recorded data, the difference between the fraction of departures from independence accounted for by the empirical and pairwise model was measured.

This captures the failures of the pairwise model that are not accounted for by the lack of reproducibility in the data. This difference was then subtracted from 100%, and the input rate r that produced an equivalent reduction was identified from the curves in Figure 9. This procedure is diagrammed with dashed lines in Figure 9 and measures the rate of common input that is consistent with the observed data.

With a common input of strength $p = 1$, an input rate of $r = 0.0034$ (equivalent to an evoked spike rate $rp/\Delta t$ in each target RGC of 0.34 Hz) produced a reduction in the fraction of departures of independence accounted for by the pairwise model equal to the value observed in the data (Fig. 9A). The data are not consistent with a common input occurring at a rate higher than r , because r subsumes any discrepancy between model and data. For comparison, the mean firing rate of the recorded ON parasol cells was 10.7 Hz. Similarly, for the pairwise-adjacent model, an input rate of $r = 0.0028$ (equivalent to an evoked RGC spike rate of 0.28 Hz) produced a value equal to the value observed in the data. As expected, for weaker common input (lower p), higher evoked rates were consistent with the values observed in the data (Fig. 9B,C). This is because lower values of p produce a much smaller proportion of multi-cell firing patterns. In summary, the data indicate that, at most, a small fraction of the observed spikes in RGCs could be produced by widely diverging common inputs in the retinal circuit.

A second test of the sensitivity of the maximum entropy approach was focused on the importance of adjacent interactions. The performance of the pairwise-adjacent model P_{adj} was compared with the performance of an alternate model. The latter was

the maximum entropy distribution subject to the same single-cell constraints $\{P(A), P(B), \dots\}$ as P_{adj} but with pairwise constraints $\{P(A, B), P(B, C), \dots\}$ obtained from a randomly selected subset of k cell pairs, where k is equal to the true number of adjacent cell pairs in the group. This pairwise-random model uses the same number of constraints as the pairwise-adjacent model but ignores the true spatial layout of recorded cells. The fraction of the departures from statistical independence accounted for by the pairwise-adjacent model, $(D_{\text{ind}} - D_{\text{adj}})/D_{\text{ind}}$, is compared with the corresponding statistic for the pairwise-random model in Figure 10, for collections of $n = 7$ ON and OFF parasol cells. The pairwise-random model exhibits substantially reduced capacity to explain the departures from independence in the data, spanning a range of performance of ~ 20 –95% rather than the values of ~ 97 –99% exhibited by the pairwise-adjacent model. As expected, a few of the values approach parity with the pairwise-adjacent model: by chance, some of the random samples will coincide with the truly adjacent pairs. These results indicate that the accuracy of the pairwise-adjacent model is not an artifact of limitations in the analysis but instead hinges critically on the true spatial layout of recorded cells, confirming that adjacent interactions are of particular importance in understanding multi-neuron firing patterns.

Discussion

Our central finding is that multi-neuron firing patterns in parasol RGCs of primate retina can be explained accurately by purely pairwise interactions restricted to adjacent cells in the mosaic. This is consistent with the simplest possible model in Figure 1 and provides a parsimonious functional description of retinal network activity. A major practical implication for future work is that large-scale visual signals conveyed from the primate retina to brain can be understood on the basis of measurements from individual cells and pairs of adjacent cells. However, the limits of this interpretation, and the implications for retinal circuitry, must be approached with caution. Below, we first discuss several implications and then return to consider the caveats.

First, the existence of multi-neuron synchrony in large collections of RGCs does not imply complex circuitry. Previous work suggested that such synchrony reflects widely diverging input from a presynaptic interneuron, such as an amacrine cell (Schnitzer and Meister, 2003). The present findings indicate that, at least in parasol cells, functional connections between pairs of adjacent cells can explain the observed synchrony. Note that this finding does not identify the mechanisms of synchrony. Previous evidence implicates a combination of mechanisms: direct gap junction coupling between neighboring RGCs (Brivanlou et al., 1998; Hu and Bloomfield, 2003; Hidaka et al., 2004; Schubert et al., 2005; Volgyi et al., 2005), gap junction coupling through intermediate amacrine cells (Dacey and Brace, 1992; Jacoby et al., 1996; Volgyi et al., 2005), and chemical synapses providing common input from bipolar or amacrine cells (Mastrorade, 1983a; Brivanlou et al., 1998) (and see DeVries, 1999; Hu and Bloomfield, 2003). The present findings do not refine this picture; instead, they reveal the circuit organization of these mechanisms in large groups of RGCs. Specifically, the following kinds of connectivity are not required to explain synchrony: (1) presynaptic common input to multiple RGCs or non-neighboring RGCs in the mosaic, or (2) reciprocal connections via intermediate cells that contact multiple or non-adjacent RGCs.

Second, the present results suggest that the spatial scale of connectivity between RGCs is smaller than the spatial scale of physiological interactions. Specifically, synchronized firing

clearly extended to pairs of cells that are not adjacent in the mosaic (Fig. 2C), but this synchrony could be explained by interactions between adjacent cells. Thus, long-range contacts, such as gap junctions at the tips of dendritic arbors (which overlap considerably in parasol cells) (Dacey and Brace, 1992) or signals propagating through wide-field amacrine cells, are not required to explain synchrony. Instead, synchrony could be caused by propagation of signals through a chain of adjacent cells in the mosaic, for example, through proximal gap junctions or narrow-field amacrine cells. Again, this finding does not uniquely identify the mechanism.

There are several caveats to the interpretations above. First, as with any model, the conclusions one may draw about retinal circuits are limited by the fact that small non-pairwise or non-adjacent interactions cannot be entirely excluded in any finite dataset. The degree to which the data quantitatively exclude such interactions are revealed by the sensitivity analysis presented in Results. Second, there are theoretical limits on what can be concluded about circuitry based on correlated firing patterns. For example, widely diverging Gaussian inputs can produce purely pairwise multi-neuron statistics. Third, only spontaneous activity and responses to a simple, fine-grained visual stimulus (white noise) were examined. It remains possible that more complex interactions, over longer distances, occur in the presence of patterned visual stimulation with more natural structure, as has been suggested in previous work (McIlwain, 1964; Olvezky et al., 2003; Schnitzer and Meister, 2003). A test of this possibility will require shuffle correcting to control for multi-cell synchrony induced by stimuli covering multiple receptive fields. Fourth, analysis was restricted to parasol cells of the primate retina, which are efficiently sampled in the present recordings. It remains possible that different cell types, and retinas of different species, exhibit more complex interactions. Finally, the present work focused on spatial patterns of activity at a single point in time. It remains possible that complex temporal patterns are introduced by interactions that are not pairwise or not adjacent. These are important avenues for future work, and the maximum entropy framework can be extended to address many of these issues.

Although maximum entropy approaches have a long history in several fields, their application in neuroscience is fairly new. Recent theoretical work has set the stage (Martignon et al., 2000; Amari, 2001; Schneidman et al., 2003b), and one group has applied the approach to recordings from salamander and guinea pig retina and cultured cortical networks, concluding that pairwise interactions explain $\sim 90\%$ of network interactions (Schneidman et al., 2003b, 2006). The present work provides several conceptual and technical advances. First, analysis was restricted to interactions between known morphological and functional types of RGC, in macaque monkey retina, and explained a substantially higher proportion of network interactions (~ 98 –99%). Second, analysis was restricted to neighboring cells in the mosaic, providing a spatial constraint on pairwise interactions within each cell type. Third, the model was cross-validated (constrained by one dataset and evaluated on another) to avoid overfitting, which can produce accurate model fits that do not generalize. Fourth, responses were measured in the presence of a steady, uniform stimulus and fine-grained white-noise stimulus rather than naturalistic stimuli or modulating uniform stimuli; the latter have coarse spatial structure and thus would be expected to produce significant departures from pairwise statistics. Finally, sensitivity analysis provided a bound on interpretation of the model in terms of network connectivity and suggested that even fairly high predic-

tion accuracy (e.g., 90%) can be consistent with substantial non-pairwise interactions in the retinal circuit.

The present findings suggest an approach to understanding the function of many large circuits in the brain, to the degree that recording technology permits. For example, the possibility of complex interactions between cells within and across columns in the neocortex (Ts'o et al., 1986; Cossart et al., 2003) or between many hippocampal neurons involved in storing spatial memories (Wilson and McNaughton, 1993) limits our capacity to understand these critical circuits. In the present work, an extremely simple pattern of connectivity sufficed to explain widespread synchrony in the retina. In some systems, such as gap junction-coupled networks in the thalamus and cortex (Chuiks Shank et al., 2005) or inferior olive (Llinas and Yarom, 1981), the present approach may translate essentially unmodified to test whether electrical coupling can fully account for network interactions. In other systems, interactions are likely to be more complicated. For example, cortical neurons may exhibit highly specific interactions dependent on layer and cell type (Yoshimura and Callaway, 2005; Yoshimura et al., 2005), resulting in complex firing patterns. The maximum entropy method is readily extended to measure the complexity of the underlying circuits. Specifically, hypothesized constraints on connectivity that emerge from anatomical considerations can be included in computation of the maximum entropy distribution, allowing a direct test of whether they account for recorded multi-cell firing patterns. For example, the approach could be used to test whether interactions are restricted to small groups of cells ($N = 2, 3, 4, \dots$), or to cells over a specific length scale such as a cortical column. The approach factors out the influence of signals propagating through intermediate cells, which otherwise would confound the analysis of group size or spatial scale. Importantly, the methods can be used with measurements of spontaneous activity in circuits such as acute slices in which natural exogenous stimulation is not possible. With the increasingly widespread use of multi-electrode recordings and optical methods *in vivo* and *in vitro* in many nervous system structures (Wilson and McNaughton, 1993; Perez-Orive et al., 2002; Beggs and Plenz, 2003; Cossart et al., 2003; Briggman et al., 2005; Ohki et al., 2005; Womelsdorf et al., 2005; Ogawa et al., 2006), the importance of understanding the complexity of network interactions has grown tremendously. Thus, extensions of the present approach may prove useful for understanding functional connectivity in other neural circuits.

References

- Amari S (2001) Information geometry on hierarchy of probability distributions. *IEEE Trans Inform Theory* 47:1701–1711.
- Arnett DW (1978) Statistical dependence between neighboring retinal ganglion cells in goldfish. *Exp Brain Res* 32:49–53.
- Beggs JM, Plenz D (2003) Neuronal avalanches in neocortical circuits. *J Neurosci* 23:11167–11177.
- Berger A, Della Pietra S, Della Pietra C (1996) A maximum entropy approach to natural language processing. *Comput Linguistics* 22:1–36.
- Briggman KL, Abarbanel HD, Kristan Jr WB (2005) Optical imaging of neuronal populations during decision-making. *Science* 307:896–901.
- Brivanlou IH, Warland DK, Meister M (1998) Mechanisms of concerted firing among retinal ganglion cells. *Neuron* 20:527–539.
- Castleman KR (1996) *Digital image processing*. Englewood Cliffs, NJ: Prentice Hall.
- Chichilnisky EJ (2001) A simple white noise analysis of neuronal light responses. *Netw Comput Neural Systems* 12:199–213.
- Chichilnisky EJ, Kalmar RS (2002) Functional asymmetries in on and off ganglion cells of primate retina. *J Neurosci* 22:2737–2747.
- Cossart R, Aronov D, Yuste R (2003) Attractor dynamics of network up states in the neocortex. *Nature* 423:283–288.
- Cover TM, Thomas JA (1991) *Elements of information theory*. New York: Wiley.
- Cruikshank SJ, Landisman CE, Mancilla JG, Connors BW (2005) Connexion connexions in the thalamocortical system. *Prog Brain Res* 149:41–57.
- Dacey DM (1993) The mosaic of midget ganglion cells in the human retina. *J Neurosci* 13:5334–5355.
- Dacey DM, Brace S (1992) A coupled network for parasol but not midget ganglion cells in the primate retina. *Vis Neurosci* 9:279–290.
- Darroch JN, Ratcliff D (1972) Generalized iterative scaling for log-linear models. *Ann Math Stat* 43:1470–1480.
- DeVries SH (1999) Correlated firing in rabbit retinal ganglion cells. *J Neurophysiol* 81:908–920.
- DeVries SH, Baylor DA (1997) Mosaic arrangement of ganglion cell receptive fields in rabbit retina. *J Neurophysiol* 78:2048–2060.
- Duda RO, Hart PE, Stork DG (2001) *Pattern classification*. New York: Wiley.
- Frechette ES, Sher A, Grivich MI, Petrusca D, Litke AM, Chichilnisky EJ (2005) Fidelity of the ensemble code for visual motion in primate retina. *J Neurophysiol* 94:119–135.
- Hertz J, Krogh A, Palmer RG (1991) *Introduction to the theory of neural computation*. Redwood City, CA: Addison-Wesley.
- Hidaka S, Akahori Y, Kurosawa Y (2004) Dendrodendritic electrical synapses between mammalian retinal ganglion cells. *J Neurosci* 24:10553–10567.
- Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA* 79:2554–2558.
- Hu EH, Bloomfield SA (2003) Gap junctional coupling underlies the short-latency spike synchrony of retinal alpha ganglion cells. *J Neurosci* 23:6768–6777.
- Jacoby RA, Stafford D, Kouyama N, Marshak D (1996) Synaptic inputs to parasol ganglion cells in the primate retina. *J Neurosci* 16:8041–8056.
- Jaynes ET (1957a) Information theory and statistical mechanics. *Physiol Rev* 106:620–630.
- Jaynes ET (1957b) Information theory and statistical mechanics II. *Physiol Rev* 108:171–190.
- Johnsen JA, Levine MW (1983) Correlation of activity in neighbouring goldfish ganglion cells: relationship between latency and lag. *J Physiol (Lond)* 345:439–449.
- Krichevsky RE, Trofimov VK (1981) The performance of universal coding. *IEEE Trans Inform Theory* 28:199–207.
- Liske J, Lemon DJ, Driver SP, Cross NJG, Couch WJ (2003) The millennium galaxy catalogue: $16 < bmgc < 24$ galaxy counts and the calibration of the local galaxy luminosity function. *Mon Not R Astron Soc* 344:307.
- Litke AM, Bezayiff N, Chichilnisky EJ, Cunningham W, Dabrowski W, Grillo AA, Grivich M, Grybos P, Hottoway P, Kachiguine S, Kalmar RS, Mathieson K, Petrusca D, Rahman M, Sher A (2004) What does the eye tell the brain? Development of a system for the large scale recording of retinal output activity. *IEEE Trans Nucl Sci* 51:1434–1440.
- Llinas R, Yarom Y (1981) Electrophysiology of mammalian inferior olivary neurones in vitro, different types of voltage-dependent ionic conductances. *J Physiol (Lond)* 315:549–567.
- Malouf R (2002) A comparison of algorithms for maximum entropy parameter estimation. In: *Proceedings of the Sixth Conference on Natural Language Learning* (Andersson B, Bergholtz M, Johannesson P, eds), pp 49–55. Stockholm: Association for Computational Linguistics.
- Marmarelis PZ, Naka K (1972) White-noise analysis of a neuron chain: an application of the wiener theory. *Science* 175:1276–1278.
- Martignon L, Deco G, Laskey K, Diamond M, Freiwald W, Vaadia E (2000) Neural coding: higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Comput* 12:2621–2653.
- Mastrorade DN (1983a) Correlated firing of cat retinal ganglion cells. I. Spontaneously active inputs to x- and y-cells. *J Neurophysiol* 49:303–324.
- Mastrorade DN (1983b) Interactions between ganglion cells in cat retina. *J Neurophysiol* 49:350–365.
- McIlwain JT (1964) Receptive fields of optic tract axons and lateral geniculate cells: peripheral extent and barbiturate sensitivity. *J Neurophysiol* 27:1154–1173.
- Meister M, Lagnado L, Baylor DA (1995) Concerted signaling by retinal ganglion cells. *Science* 270:1207–1210.
- Nemenman I, Shafee F, Bialek W (2002) Entropy and inference, revisited.

- In: *Advances in neural information processing systems*, Vol 14 (Dietterich TG, Becker S, Ghahramani Z, eds). Cambridge, MA: MIT.
- Nirenberg S, Carcieri SM, Jacobs AL, Latham PE (2001) Retinal ganglion cells act largely as independent encoders. *Nature* 411:698–701.
- Ogawa H, Cummins GI, Jacobs GA, P J Miller (2006) Visualization of ensemble activity patterns of mechanosensory afferents in the cricket cercal sensory system with calcium imaging. *J Neurobiol* 66:293–307.
- Ohki K, Chung S, Ch'ng YH, Kara P, Reid RC (2005) Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature* 433:597–603.
- Olveczky BP, Baccus SA, Meister M (2003) Segregation of object and background motion in the retina. *Nature* 423:401–408.
- Orlitsky A, Santhanam N, Zhang J (2004) Universal compression of memoryless sources over unknown alphabets. *IEEE Trans Inform Theory* 50:1469–1481.
- Paninski L (2003) Estimation of entropy and mutual information. *Neural Comput* 15:1191–1253.
- Peichl L, Wässle H (1981) Morphological identification of on- and off-centre brisk transient (γ) cells in the cat retina. *Proc R Soc Lond B Biol Sci* 212:139–153.
- Perez-Orive J, Mazor O, Turner GC, Cassenaer S, Wilson RI, Laurent G (2002) Oscillations and sparsening of odor representations in the mushroom body. *Science* 297:359–365.
- Perkel DH, Gerstein GL, Smith MS, Tatton WG (1975) Nerve-impulse patterns: a quantitative display technique for three neurons. *Brain Res* 100:271–296.
- Polyak SL (1941) *The retina*. Chicago: University of Chicago.
- Press WH, Flannery BP, Teukolsky SA, Vetterling WT (1988) *Numerical recipes in C*. Cambridge, UK: Cambridge UP.
- Rodieke RW (1998) *The first steps in seeing*. Sunderland, MA: Sinauer.
- Roerdink J, Meijster A (2000) The watershed transform: definitions, algorithms and parallelization strategies. *Fundamenta Informaticae* 41:187–228.
- Schneidman E, Bialek W, Berry II MJ (2003a) Synergy, redundancy, and independence in population codes. *J Neurosci* 23:11539–11553.
- Schneidman E, Still S, Berry II MJ, Bialek W (2003b) Network information and connected correlations. *Phys Rev Lett* 91:238701.
- Schneidman E, Berry MJ, Segev R, Bialek W (2006) Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440:1007–1012.
- Schnitzer MJ, Meister M (2003) Multineuronal firing patterns in the signal from eye to brain. *Neuron* 37:499–511.
- Schubert T, Degen J, Willecke K, Hormuzdi SG, Monyer H, Weiler R (2005) Connexin36 mediates gap junctional coupling of alpha-ganglion cells in mouse retina. *J Comp Neurol* 485:191–201.
- Segev R, Goodhouse J, Puchalla J, Berry MJ (2004) Recording spikes from a large fraction of the ganglion cells in a retinal patch. *Nat Neurosci* 7:1154–1161.
- Ts'o DY, Gilbert CD, Wiesel TN (1986) Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J Neurosci* 6:1160–1170.
- Volgyi B, Abrams J, Paul DL, Bloomfield SA (2005) Morphology and tracer coupling pattern of alpha ganglion cells in the mouse retina. *J Comp Neurol* 492:66–77.
- Wandell BA (1995) *Foundations of vision*. Sunderland, MA: Sinauer.
- Wässle H, Peichl L, Boycott BB (1983) Mosaics and territories of cat retinal ganglion cells. *Prog Brain Res* 58:183–190.
- Watanabe M, Rodieke RW (1989) Parasol and midget ganglion cells of the primate retina. *J Comp Neurol* 289:434–454.
- Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261:1055–1058.
- Womelsdorf T, Fries P, Mitra PP, Desimone R (2005) Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* 439:733–736.
- Yoshimura Y, Callaway EM (2005) Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity. *Nat Neurosci* 8:1552–1559.
- Yoshimura Y, Dantzker JL, Callaway EM (2005) Excitatory cortical neurons form fine-scale functional networks. *Nature* 433:868–873.