Biol Cybern (2006) 95:327–348
DOI 10.1007/s00422-006-0093-5

ORIGINAL PAPER

Biological Cybernetics

See globally, spike locally: oscillations in a retinal model encode large visual features

Greg J. Stephens · Sergio Neuenschwander · John S. George · Wolf Singer · Garrett T. Kenyon

Received: 25 May 2005 / Accepted: 29 May 2006 / Published online: 9 August 2006

Abstract We show that coherent oscillations among neighboring ganglion cells in a retinal model encode global topological properties, such as size, that cannot be deduced unambiguously from their local, time-averaged firing rates. Whereas ganglion cells may fire similar numbers of spikes in response to both small and large spots, only large spots evoke coherent high frequency oscillations, potentially allowing downstream neurons to infer global stimulus properties from their local afferents. To determine whether such information might be extracted over physiologically realistic spatial and temporal scales, we analyzed artificial spike trains whose oscillatory correlations were similar to those measured experimentally. Oscillatory power in the upper gamma band, extracted on single-trials from multi-unit spike trains, supported good to excellent size discrimination between small and large spots, with performance improving as the number of cells and/or duration of the analysis window was increased. By using Poisson distributed spikes to normalize the firing rate across stimulus conditions, we further found that coincidence detection, or synchrony, yielded substantially poorer performance on identical size discrimination tasks. To determine whether size encoding depended on contiguity independent of object shape, we examined the total oscillatory activity across the entire model retina in response to random binary images. As the ON-pixel probability crossed the percolation threshold, which marks the sudden emergence of large connected clusters, the total gamma-band activity exhibited a sharp transition, a phenomena that may be experimentally observable. Finally, a reanalysis of previously published oscillatory responses from cat ganglion cells revealed size encoding consistent with that predicted by the retinal model.

G. J. Stephens · J. S. George · G. T. Kenyon
Physics Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

G. J. Stephens
Lewis-Sigler Institute for Integrative Genomics, Carl Icahn Laboratory, Princeton University, Princeton, NJ 08544, USA

S. Neuenschwander · W. Singer
Max-Planck-Institut für Hirnforschung, Deutschordenstraße 46, 60528 Frankfurt am Main, Germany

G. T. Kenyon (✉)
Applied Modern Physics, MS D454, Los Alamos National Laboratory, Los Alamos, NM 87545, USA
e-mail: gkenyon@lanl.gov

1 Introduction

High frequency oscillations have been observed in a variety of vertebrate retinas, including cat (Laufer and Verzeano 1967; Neuenschwander et al. 1999; Neuenschwander and Singer 1996; Steinberg 1966), rabbit (Ariel et al. 1983), mudpuppy (Wachtmeister and Dowling 1978), frog (Ishikane et al. 2005, 1999), macaque (Frishman et al. 2000) and human (De Carli et al. 2001; Wachtmeister 1998). Nonetheless, the role of coherent high frequency oscillations in visual processing remains unresolved. Here, we used computer models to investigate whether coherent oscillations among neighboring retinal ganglion cells could encode global stimulus properties, such as size, that cannot be unambiguously inferred from their local, time-averaged firing rates. Ganglion cell firing rates typically depend both on luminance contrast and on stimulus size, increasing for spot

diameters smaller than the receptive field center and decreasing thereafter due to surround inhibition. Because it is impossible, based entirely on local information, to disentangle the effects of global stimulus topology from the effects of luminance contrast, the size of a stimulus cannot be unambiguously determined from the mean firing rates of one or a few neighboring cells. Coherent high frequency oscillations, on the other hand, are not evoked by small spots (Ariel et al. 1983; Ishikane et al. 1999; Neuenschwander et al. 1999) and thus are not intrinsically confounded by luminance contrast.

We propose that the retina encodes visual information in two complementary channels: while the average firing rate among a small group of neighboring ganglion cells conveys local stimulus properties, such as luminance contrast, their coherent oscillations convey additional information about global stimulus properties, such as size. To test this hypothesis, we asked whether global topological information could be extracted from the coherent oscillations among neighboring retinal ganglion cells upon single stimulus presentations lasting no more than a few hundred msec, equivalent to typical inter-saccade intervals (Martinez-Conde et al. 2004). The answer to this question both constrains and informs the functional role of retinal oscillations. In order to be behaviorally relevant, any information conveyed by coherent oscillations must be accessible on physiological time scales. Furthermore, spatial convergence at early visual processing stages is rather low, being one to a few at retina-LGN synapses (Usrey et al. 1999) and approximately 1:15 at LGN synapses onto layer IV simple cells if ON and OFF inputs are considered separately (Alonso et al. 2001). If global stimulus properties cannot be extracted from a relatively small number of neighboring ganglion cells on single trials, then it is unlikely that coherent oscillations convey such information to early processing stages. Alternatively, by demonstrating that retinal oscillations encode global topological information that can, in principle, be extracted early in the visual processing hierarchy, we establish their potential functional utility.

To investigate whether global stimulus properties could be extracted from the coherent oscillations among relatively small numbers of retinal ganglion cells on single trials, we analyzed sets of artificial spike trains produced by a computer model of the inner retina (Kenyon et al. 2003). Coherent oscillations are an emergent property of large neural ensembles and thus can be difficult to study using standard electrophysiological techniques. Computational models can therefore provide a useful alternative for investigating the information encoded by coherent oscillations among specific cell types. The retinal model produced size-dependent coherent oscillations similar to those measured experimentally, as assessed by both single-trial as well as trial-averaged multi-unit correlation functions and corresponding frequency spectra. A discrimination analysis was used to classify the relative sizes of spot stimuli as either "smaller" or "larger" on the basis of the oscillations present in short sections of multi-unit spike train data, from 50 to 400 m s in duration and containing between 1 and 16 model ganglion cells. Despite the presence of substantial single-trial variability, it was nonetheless possible to discriminate between large and small spots on between 60 and 100% of stimulus trials, with performance improving as the duration of the analysis window was increased and/or as more neurons were included in the multi-unit spike train. Thus, our results suggest that information in the optic nerve may be multiplexed, with local stimulus properties, such as contrast, encoded by the firing rates of individual neurons and contextual information, such as whether a given ganglion cell is responding to a small isolated spot or to a large connected object, encoded by coherent high frequency oscillations.

To control for size-dependent changes in the mean firing rate, the above discrimination analysis was repeated after adding Poisson distributed events to the individual multi-unit records in order to normalize the average number of spikes produced under different stimulus conditions. Eliminating all size information from the mean firing rate did not qualitatively reduce the ability to distinguish between large and small spots based on the strength of the coherent oscillations between neighboring ganglion cells. In fact, by reducing the baseline variability in the responses to small spots, rate normalization actually led to slightly improved performance on the size discrimination task. Thus, our results suggest that coherent oscillations can encode global topological information on single trials even when the average power in the DC band is held constant. Employing the same rate normalization technique, we found that size discrimination based on coincidence detection, or synchrony, was substantially worse than that mediated by coherent oscillations, even though both tasks used the same multi-unit spike records. Our results thus suggest that single trial power spectra may convey more information about global stimulus properties, such as size, than does the absolute number of coincident inputs, although our results say nothing about other information processing roles for firing synchrony that have been suggested (Meister and Berry 1999; Schnitzer and Meister 2003; Singer and Gray 1995).

We further investigated whether our results were critically dependent on the details of the retinal model, particularly the axon-mediated feedback used to reproduce

the high frequency oscillations recorded from cat ganglion cells. Simply because the retinal model was able to reproduce the trial-averaged correlations observed experimentally does not necessarily mean that the single-trial statistics were correct. For example, the signal-to-noise in the retinal model could have been substantially different from that present in the retina itself and yet this discrepancy might have been obscured in the process of averaging over multiple trials. We therefore generated a second set of artificial spike trains that employed a very different mechanism for producing realistic high frequency oscillations, namely, a common modulatory input. Both sets of artificial spikes trains supported approximately the same levels of performance on the size discrimination task, suggesting that the contextual information encoded by coherent high frequency oscillations may be largely independent of their mode of generation.

In addition, we investigated the degree to which the retinal oscillations evoked by large contiguous objects were independent of shape. Random binary images were used to generate different sized clusters with highly variable boundaries. Rather than analyzing the responses to each cluster individually, we examined the output across the entire retinal model for an analog of the "percolation transition" in the set of input images. As the average density of ON (or equivalently, OFF) pixels in a set of random binary images approaches a critical value, called the percolation threshold (Grimmett 1999), the average size of a cluster of simply connected ON (or OFF) pixels diverges, becoming commensurate with the scale of the image itself. By measuring the total activity in the upper gamma band across the entire model retina, we were able to test the hypothesis that high frequency oscillations could be evoked by large connected objects of arbitrary shape. Consistent with this hypothesis, there was a sharp transition in the total gamma activity near the percolation threshold at which large connected clusters first appear. These results suggest a novel experimental strategy for testing some of the key predictions of the retinal model.

Finally, we reanalyzed previously published multiunit recordings from cat retinal ganglion cells that were stimulated by several different spot sizes. Because the experimental study did not identify individual cells, it was not possible to determine what fraction of the recorded neurons participated in the high frequency oscillatory responses, nor was it possible to determine whether the multi-unit data arose from a fixed set of neurons or, alternatively, if separate cell populations were activated by the different sized stimuli. In addition, the mean firing rate declined by approximately a factor of two over the course of the response, which may have further obscured the information encoded by coherent oscillations. Despite such confounds, our reanalysis was clearly consistent with the hypothesis that coherent high frequency oscillations between neighboring retinal ganglion cells can provide fast, local encoding of global stimulus properties.

Preliminary descriptions of the results reported here have appeared previously in abstract form (Stephens et al. 2003).

2 Methods

2.1 Retinal model

Artificial spike trains were generated using a semi-realistic model of the inner retina whose relationship to anatomical and physiological data (Kenyon et al. 2003), ability to accomplish rudimentary segmentations of both artificial and naturalistic scenes (Kenyon et al. 2004a, c) and ability to compensate for correlation-induced limits on rate coded stimulus features, such as intensity (Kenyon et al. 2004b), have been previously described. Briefly, input to the model was conveyed by an array of external currents proportional to the pixel-by-pixel grayscale value of a two-dimensional image. These external currents directly stimulated the model bipolar cells and approximated their light-modulated synaptic input from cone photoreceptors. The bipolar cells produced excitatory postsynaptic potentials in both ganglion cells and amacrine cells according to a random process (Freed 2000). The axon-bearing amacrine cells were electrically coupled to neighboring ganglion cells and to each other and made strong inhibitory connections onto the surrounding ganglion cells and axon-bearing amacrine cells (Fig. 1). This feedback circuit produced robust, physiologically realistic oscillations in response to large stimuli. When several ganglion cells were activated by a stimulus, they in turn activated neighboring axon-bearing amacrine cells via gap junctions (Dacey and Brace 1992; Jacoby et al. 1996; Vaney 1994). The stimulated cells were then hyperpolarized by the ensuing wave of axon-mediated inhibition, thus setting up the next cycle of the oscillation. Spike generation was modeled as a leaky integrate-and-fire process with a membrane time constant of 5 m s, consistent with published physiological data from cat alpha ganglion cells (O'Brien et al. 2002). The model also contained local non-spiking amacrine cells that generated randomly distributed inhibitory postsynaptic potentials that helped both to make spontaneous firing asynchronous and to increase the overall dynamic range, but were not otherwise critical for generating oscillatory spike trains.

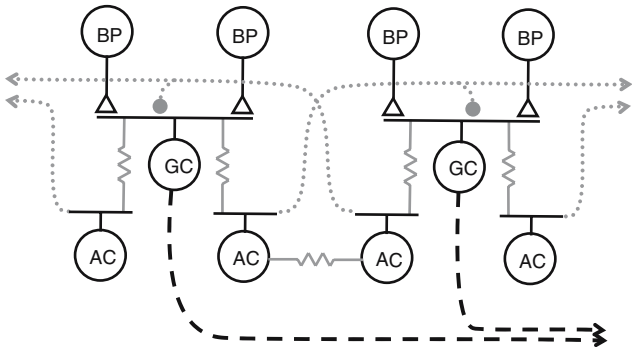

Fig. 1 Schematic of the retinal feedback circuit. Only those elements directly responsible for synchronous oscillations are depicted. A combination of local excitation via gap junctions (resistors) and long-range inhibition via axon-bearing amacrine cells (*gray dotted lines* and *filled black circles*) produced physiologically realistic size-dependent oscillations. The entire model contained a 32×32 array of ganglion cells driven by a 64×64 array of bipolar cells, with inhibitory feedback arising from a 64×64 array of axon-bearing amacrine cells, only a few of which are shown. Light stimuli were implemented by injecting currents directly into bipolar cells

Finally, Poisson distributed events were added to the spike trains generated by the network model so as to increase the mean firing rate by 20% regardless of stimulus size (these random events were added in lieu of, not in combination with, the rate normalization process described below). This procedure yielded correlations that better matched published trial-averaged experimental data, possibly because multi-unit spike trains recorded from the cat retina may have included some non-oscillatory cells.

A formal mathematical description of model is presented in Appendix A. The robustness of the simulations with respect to both numerical and physiological parameters (Kenyon et al. 2003) as well as an earlier linear version of the model (Kenyon and Marshak 1998), have been presented previously.

2.2 Common input model

A second set of artificial spike trains were generated under the alternative assumption that the spatiotemporal correlations between cat ganglion cells are due to a common oscillatory input that uniformly modulates the firing rates of all simultaneously recorded neurons. Starting with a power spectrum containing a single Gaussian peak, characterized by three parameters specifying the amplitude, width, and central frequency, we transformed back to the time domain by randomly choosing the phases of the individual Fourier components. Parallel sets of artificial spike trains with realistic spatial–temporal correlations were then constructed by assuming an array of Poisson generators, where the above time series was used to simultaneously modulate their instantaneous firing rates. Further mathematical details of the procedure are presented in Appendix B as well as in previous studies (Kenyon et al. 2004b).

2.3 Correlation analysis

Correlations in the artificially generated spike trains were computed from 200 m s epochs, which in the case of the retinal circuit model were taken from the sustained portion of the response after the firing activity had settled to a stable plateau level. Correlations were computed separately between all distinct cell pairs contributing to the multi-unit record and the individual cross-correlations then averaged to produce a single multi-unit correlation function. All multi-unit correlation functions were normalized relative to the expected correlations due to chance. An amplitude of one thus corresponded to a doubling in the number of correlated events over the expected value at the corresponding delay. The single-trial correlation functions were averaged across all 200 independent stimulus trials to produce a trial-averaged multi-unit correlation function. All correlation functions used a bin width of 1 m s and were lag corrected for edge effects arising from the finite length of the spike train. Shift-predictors, which were generally negligible, were not subtracted, as our hypothesis does not depend on whether oscillations are phase-locked to the stimulus onset.

2.4 Size discrimination

We assessed the ability to discriminate between small and large stimuli based on the average spectral amplitude in the upper gamma frequency band (either 65–100, 70–90 or 60–95 Hz, as indicated), extracted from short segments of multi-unit spike train data on single trials. Model spike train segments were 50–400 ms in duration, contained from 1–16 neurons and for the retinal circuit model were always taken from the plateau portion of the response, beginning at least 200 m s after stimulus onset. To estimate the average spectral amplitude in the upper gamma band, the multi-unit spike train was Fourier transformed and the mean amplitude of the discrete Fourier coefficients between the indicated limits was computed. The individual Fourier coefficients were then scaled by the total number of spikes, given by the amplitude of the zero-frequency band, thus removing the linear dependence on the mean firing rate. For spike trains that were not rate-normalized (the rate normalization procedure is described below), the average spectral amplitude in the upper gamma band was expressed

as a fraction of the mean baseline spectral amplitude, defined as the average between 220 and 500 Hz, computed on the same trial. This alternative scaling procedure was necessary because the average power in the upper gamma band exhibited a non-linear dependence on the total number of spikes. Trial-to-trial fluctuations became larger as the number of spikes decreased but the amplitude could not drop below zero, thus creating a net positive baseline shift when the firing rate was very low. The average spectral amplitude in the upper gamma band (gamma activity) was determined for each trial and the results sorted into 11 equally spaced bins spanning the full range of the data. The binned, single-trial gamma activity was then normalized as a probability distribution, yielding the percentage of trials on which the single-trial gamma activity fell within a particular range.

Given the probability distributions of single-trial gamma activity for each stimulus size, it was straightforward to estimate the percentage of trials on which any given pair of stimulus sizes could be discriminated based on the single-trial gamma activity alone. Specifically, the ability to discriminate two stimulus sizes was inversely related to the degree of overlap between the corresponding probability distributions (Duda et al. 2001). If the distributions of single-trial gamma activity overlapped completely, the maximal theoretical performance on the size discrimination task would be no better than chance (50% correct). On the other hand, if the distributions of single-trial gamma activity were entirely non-overlapping for a given pair of stimulus sizes, the maximum theoretical performance on the discrimination task would be perfect (100% correct). Between these two extremes, corresponding to distributions that partially overlap, maximum theoretical performance on the size discrimination task, P, expressed as a fraction of trials correctly classified, is given by the following formula:

$$P = \frac{2 - A_{\text{overlap}}}{2} \tag{1}$$

where A_{overlap} denotes the total area of the overlap between the two distributions and the maximum value of A_{overlap} is normalized to one. Error bars on the estimated values of P were determined by assuming the number of trials to either side of the Bayes discriminator obeyed binomial statistics.

2.5 Rate normalization

To estimate the size information conveyed by the number of synchronous events on each trial, as opposed to the information conveyed by coherent oscillations, it was useful to first normalize the average firing rates across stimulus conditions so that the baseline level of coincident inputs remained constant. For the common input model, the average firing rate was fixed a priori at 50 Hz for all spot sizes, so explicit rate normalization was unnecessary. For the retinal circuit model, the average firing rate depended on spot size, primarily because the small spots did not completely cover the receptive field centers of the recorded cells. To normalize the mean response, we added independently generated Poisson distributed spikes to each train so that the average firing rate was always equal to 50 Hz independent of stimulus size. A similar procedure was used to normalize the firing rates of the spike trains recorded experimentally, as described below.

2.6 Physiological recordings

Data from previously reported experiments in the cat retina were reanalyzed in order to quantify the information relating to the global stimulus topology conveyed by high frequency oscillations in short sections of multi-unit spike train data. As no new experiments were conducted, and since a detailed description of the experimental methods is available elsewhere (Neuenschwander et al. 1999), only an abbreviated description of the experimental procedure is provided here.

Intra-ocular recordings were made from anesthetized and paralyzed cats in response to spots of various sizes, presented at high contrast against a background illumination of $\sim 0.4\,\text{cdm}^{-2}$. Each spot of a given size was presented 20 times for a period of 3 s, although only the first 2 s of the response were used in the following analysis. For each spot size, the total spike record across all trials was divided into 200 equally spaced, non-overlapping segments, each 200 m s in length. For purposes of this study, these 200 spike train segments were treated as independent stimulus trials. Because single units were not isolated in the experimental data, correlations were computed directly from the multi-unit spike train rather than between all distinct cell pairs. Thus, unlike the correlations computed from artificially generated spike trains, the correlations computed from the cat data included the auto-correlation functions of the individual cells. The rate normalization procedure, when used, was modified so as to equalize the average firing rates in each 200 m s epoch relative to the stimulus onset, as well as with respect to spot size. In all other aspects, the experimental and artificial spike trains were processed identically.

3 Results

3.1 Analysis of representative single-trials

The retinal circuit model (Fig. 1) was stimulated using square spots covering an area equal to either 1×1, 4×4, or 6×6 ganglion cells (sizes expressed as multiples of the ganglion cell center-to-center distance, corresponding to approximately 1°–4° degrees of visual angle, intensity was equal to 0.25 for all stimuli). Representative multi-unit event trains, consisting of four model neurons taken from a 2×2 array of neighboring ganglion cells located at the center of the stimulus, were obtained in response to three different spot sizes (Fig. 2a). Spikes generated by the retinal circuit model, including the addition of randomly distributed Poisson events that increased the average firing rate by 20%, are shown in black. Rate normalized spike trains, in which Poisson distributed events were added to maintain the average firing rate at 50 Hz regardless of spot size, are shown overlaid with the additional spikes drawn in light gray. Because the smallest spot did not completely cover the receptive field centers of the four recorded neurons, the average number of spikes varied by approximately a factor of 3 as a function of stimulus size (black spikes only). Based purely upon visual examination of the rate-normalized event trains (gray + black spikes), it is difficult to distinguish between the different spot sizes, although there is some suggestion of greater periodicity in the responses to the larger stimuli. The primary question we address in this study is whether there is enough periodic structure in the single-trial, multi-unit data recorded from a few neighboring ganglion cells, even in the rate-normalized case, to infer global stimulus properties, such as size.

Single-trial frequency spectra computed from a set of representative non-rate-normalized multi-unit spike trains (i.e. without gray spikes) reveal that as a function of stimulus size there is a prominent increase, relative to the surrounding baseline, in spectral amplitudes falling within the upper gamma band, roughly between 60 and 100 Hz (Fig. 3a). The magnitudes of the Fourier coefficients computed from individual spike train segments were divided by the total number of events, given by the amplitude at 0 Hz, or DC band. In response to the largest spot, covering an array of 6×6 model ganglion cells, all three single-trial traces rose to a distinct peak at approximately 85 Hz, although the exact location varied from trial-to-trial. On the other hand, in response to the smallest spot, whose diameter was equal to only 1 ganglion cell center-to-center distance, there was no evidence of a distinct spectral peak anywhere within

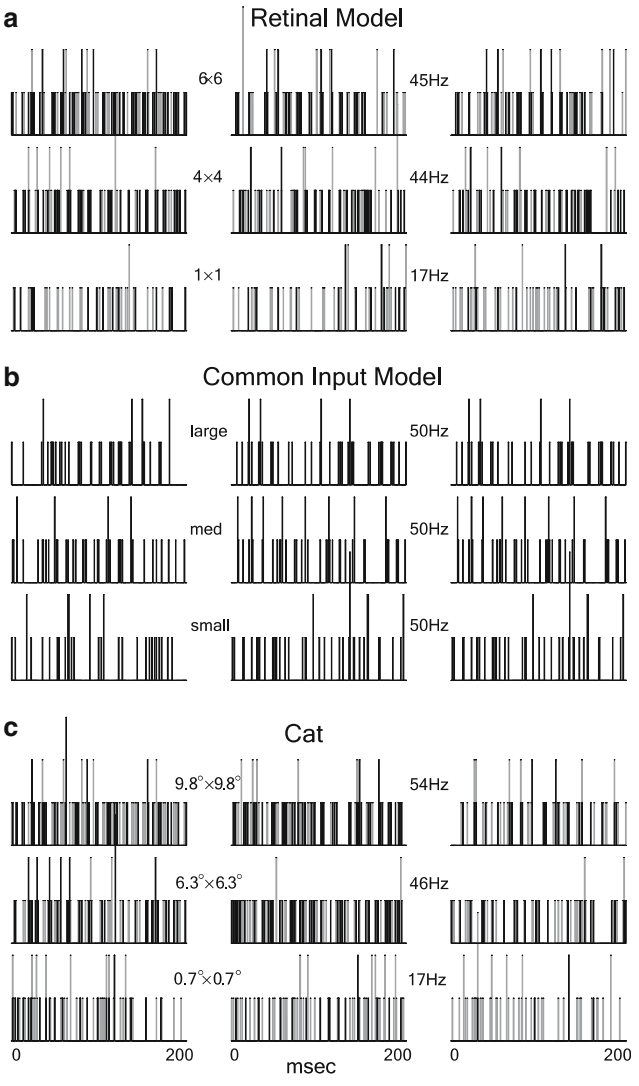

Fig. 2 Representative multi-unit spikes trains. *Gray bars* indicate Poisson distributed events added to maintain the average firing rate constant across stimulus conditions (trial-averaged firing rate of the original spikes trains indicated to upper right of center). Spike trains were 200 m s in duration and stimuli were maintained throughout the response. **a**–**c** *Top row* Three representative responses to the largest applied stimulus (spot dimensions indicated to upper left of center). *Middle row* Responses to an intermediate sized spot. *Bottom row* Responses to the smallest spot. Upon visual inspection, the rate-normalized spike trains exhibit only subtle periodic structure as a function of increasing stimulus size. **a** Retinal circuit model. Multi-unit spike trains recorded from 2×2 array of model ganglion cells centered on the applied stimulus. Data taken from the sustained portion of the response after the firing rate had settled to a stable plateau level. **b** Common input model. Each sequence contained 4 units. Parameters corresponding to large, med and small spots were chosen to provide a reasonable match to experimentally measured size-dependent correlations (see Fig. 6). **c** Cat. From left to right, spike trains recorded 0.2, 1.0 and 2.0 s after stimulus onset. The mean firing rate adapted over this period by approximately a factor of two

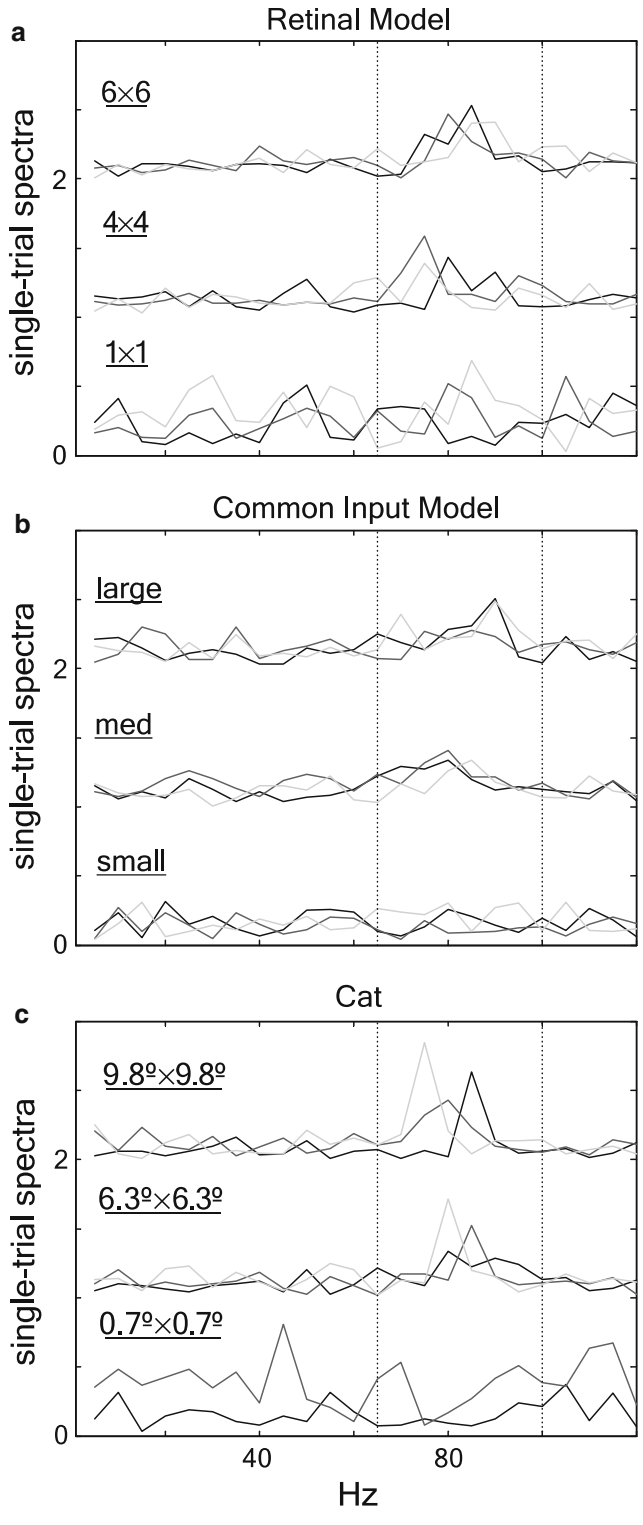

◄ **Fig. 3** Representative frequency spectra as a function of spot size. Each trace was computed from a single non-rate-normalized spike train, examples of which were shown in Fig. 2 (black events only). Spot dimensions are indicated to upper left of each trace. Spectral amplitudes expressed relative to the number of spikes (height of DC band). **a** Retinal circuit model. For the small spot, increased variability was due to the reduced number of events. **b** Common input model. **c** Cat. In the responses to the small spot, variability was larger for traces recorded at times later from stimulus onset. *Black, dark and light gray* traces recorded 0.2, 1.0 and 2.0 s following stimulus onset, respectively. **a–c** An increase in spectral amplitudes in the upper gamma band is clearly evident as a function of increasing spot size, suggesting that coherent oscillations can encode topological information that is available on behaviorally relevant time scales. *Vertical dotted lines* denote the range 65–100 Hz

the upper gamma band. Single-trail frequency spectra obtained in response to the small spot were noisier due to the reduced number of spikes in a typical sequence. These results suggest that by analyzing the gamma activity present in short sections of multi-unit data containing only a few neighboring ganglion cells, it may be possible to reliably infer global stimulus properties over physiologically relevant time scales.

The size-dependent peaks in the single-trial frequency spectra recorded from the retinal circuit model were evident in the trial-averaged responses as well (Fig. 4a). For the largest spot, the peak amplitude in the trial-averaged data was slightly more than double the average baseline level, whether computed using the original (black traces) or the rate-normalized spike trains (gray traces). Importantly, there was no evidence of a peak in the upper gamma band in response to the smallest stimulus, although the rate normalization procedure did reduce the baseline offset, a non-linear effect that may be attributed to the reduced variability in the single-trial traces that diminished the influence of the amplitude cutoff at zero. The rate normalization procedure did not much affect the mean spectral amplitudes computed in response to the two larger stimuli, as only relatively few Poisson distributed spikes were added under these conditions. However, there did appear to be a general reduction of total gamma activity in the trial-averaged responses compared to the single trial traces, most likely due to the intrinsic variability in the location of the peak oscillation frequency, itself a consequence of the non-linear nature of the retinal feedback circuit (Kenyon et al. 2003).

Similar single-trial size discrimination was evident in the multi-unit correlation functions computed from representative spike sequences (Fig. 5a). Rate-normalized spike trains were used for this analysis since it was impractical to obtain single-trial correlation functions when the total number of events was very small. Although the individual traces were quite noisy, periodic structure was still clearly evident in the responses to the large spot, whereas no clear gamma activity was evi-

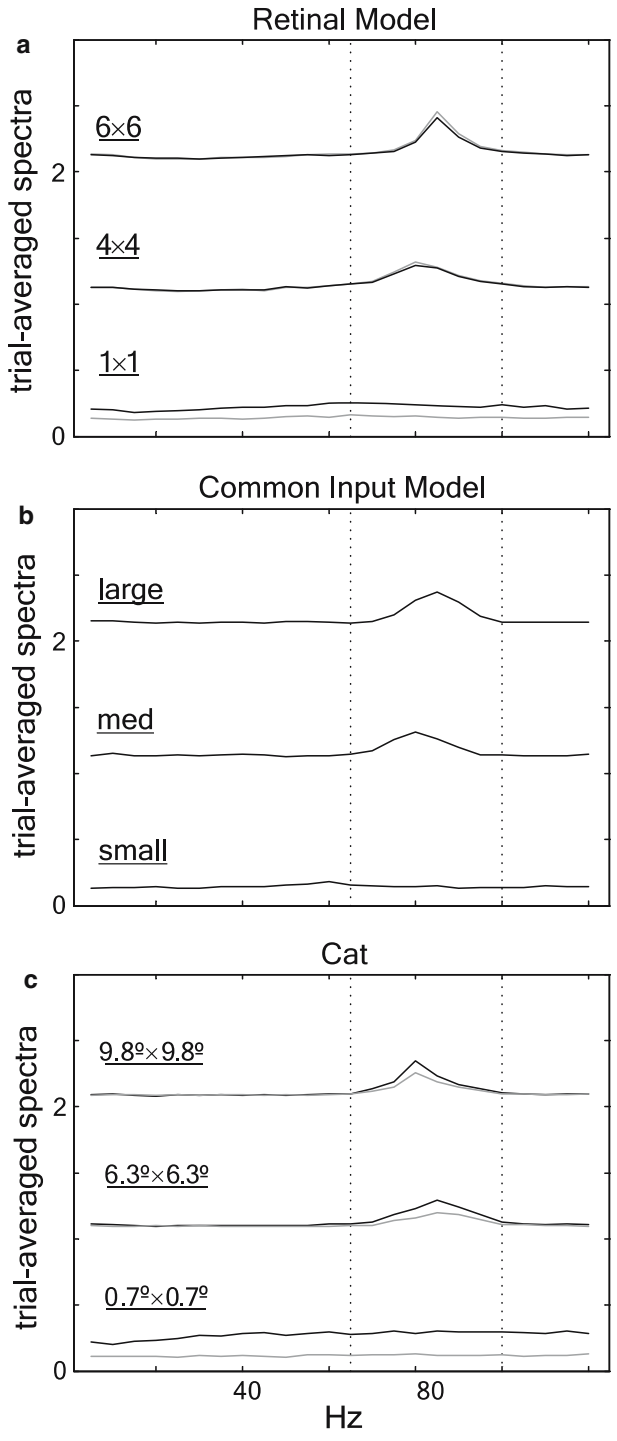

Fig. 4 Trial-averaged frequency spectra as a function of spot size. Same general organization as in Fig. 3. **a–c** A prominent peak in the upper gamma band appears as spot size increases. *Black traces* were computed using the original spike trains. *Gray traces* were computed using the rate-normalized data. For both the retinal circuit model and the cat data, the rate normalization procedure produced a general reduction in all frequency components obtained in response to the small spot but did not affect the overall shape of the trial-average spectra. Rate-normalization also reduced the size-dependent peaks computed from the cat data, by diluting gamma activity relative to the DC band

dent in the single-trial correlations functions computed in response the small spot. Consistent with the trend suggested by the single-trial traces, the trial-averaged multi-unit correlations functions were also strongly size dependent (Fig. 6a), an effect that has been observed in several species (Ariel et al. 1983; Ishikane et al. 1999; Neuenschwander et al. 1999). However, the averaging process seemed to reduce the magnitude of the oscillations from that present in the single trial traces, particularly the height of the periodic side bands, an effect that may be attributed to the slight misalignment of corresponding maxima across trials. Despite such variations, the strong size dependence visible in the single-trial correlation functions suggest that global topological information is likely to be available on behaviorally significant spatial and temporal scales, in this case represented by spike train segments 200 m s in duration and consisting of four neighboring cells.

Analogous results were obtained using an alternative model based on common modulatory input, a method for generating oscillatory spike trains that is fundamentally different from, and largely complementary to, the axon-mediated feedback employed in the retinal circuit model. Upon casual inspection, representative multi-unit spike trains generated by the common input model exhibited only subtle differences in periodic structure as a function of spot size, even though the instantaneous firing rate was strongly modulated by physiologically realistic coherent oscillations (Fig. 2b). As with the retinal circuit model, the artificial spike trains generated by the common input model contained four units. The spike trains corresponding to the different stimulus sizes have been labelled “small”, “med” and “large” to facilitate comparison. By adjusting the amplitudes and widths of the gamma activity peaks in the corresponding frequency spectra and then transforming back to the time domain after assigning random phases to each frequency component (see Methods), it was possible to produce a reasonable fit to the size dependent correlations measured experimentally. Because the time-averaged firing rate in the common input model was held constant at 50 Hz across all stimulus conditions, there was no need for an additional rate normalization step (i.e. no gray spikes).

As in the retinal circuit model, the single-trial frequency spectra produced by common modulatory input showed a clear increase in spectral amplitudes falling within the upper gamma band as a function of increasing stimulus size (Fig. 3b), an effect that was present in the trial-averaged responses as well (Fig. 4b). One difference between the two models, however, was that the total gamma activity appeared to be somewhat smaller

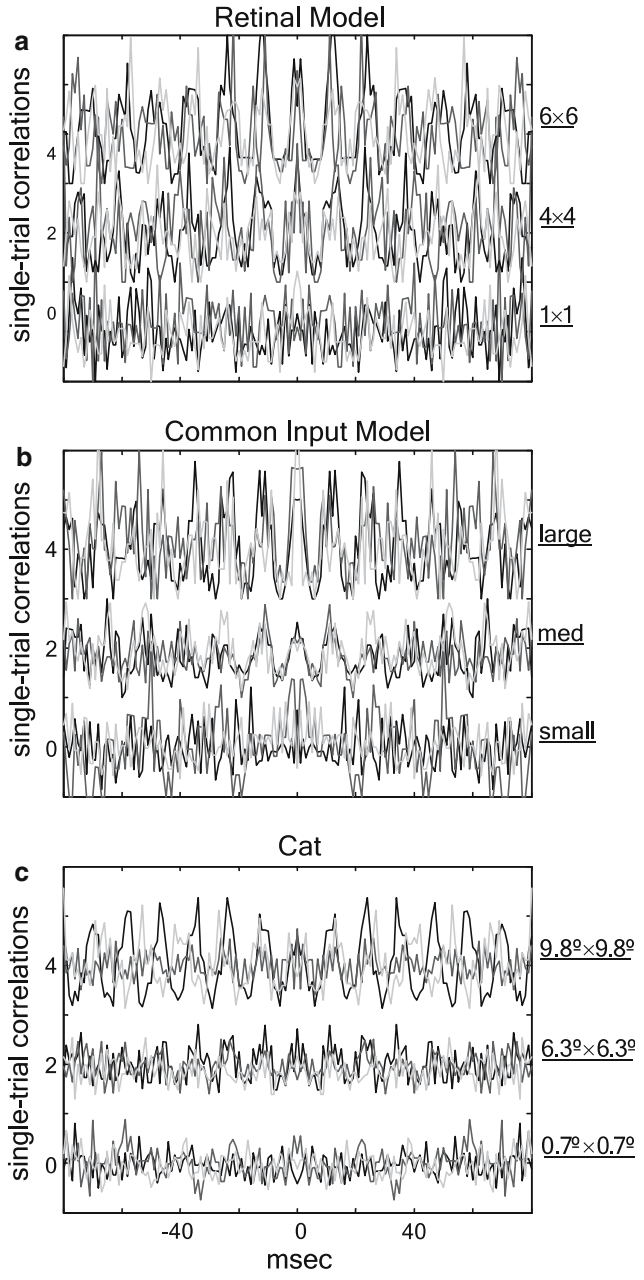

Fig. 5 Single-trial correlation functions for different sized spots. Same general organization as in Fig. 3. **a–c** Each trace computed from a single rate-normalized spike train (e.g. black + gray events in Fig. 2). Correlations were expressed as a fraction of the expected level due to chance. Increases in oscillatory structure were clearly evident as a function of increasing stimulus size. Coherent oscillations evoked by the largest spot size could persist over many tens of m s, although corresponding maxima were often poorly aligned across trials

in the single trial traces produced by common input, as opposed to the retinal feedback circuit, although the areas under their trial-averaged peaks were similar. A possible explanation of this difference is that the trial-to-trial variability in the location of the peak frequency may have been greater in the retinal circuit model, as the location of the peak frequency in the common input model was held fixed for each stimulus condition. The two models were nonetheless similar in that the single-trial correlation functions produced by common oscillatory input (Fig. 5b) again appeared to exhibit larger periodic modulations than were present in the trial-averaged results (Fig. 6b). Although the central oscillation frequency was held fixed in the common input model, the instantaneous phase of the underlying coherent oscillations varied randomly over time, so that the locations of the *n*th-order side peaks were often poorly aligned across trials. Overall, the results from the common input model both confirm the primary predictions of the retinal circuit model and suggest that the global topological information conveyed by coherent oscillations between small groups of neighboring neurons may be largely independent of their mode of generation.

For comparison, we also considered representative multi-unit spike trains recorded from the cat retina in response to several different stimulus sizes (Neuenschwander et al. 1999), where each sequence consisted of approximately 4–5 ganglion cells (Fig. 1c, black spikes). A rate normalization procedure, analogous to that applied to the spike trains generated by the retinal circuit model, was used to isolate the information encoded by coherent oscillations independent of changes in the average firing rate (added spikes shown in gray). Due to adaptation effects, the average firing rate fell monotonically throughout the stimulus. Representative multi-unit spike trains were chosen to illustrate the different phases of the response, beginning either 0.2, 1.2 and 2.0 s following stimulus onset (left, middle and right columns, respectively). As was the case with the artificially generated spike trains, the rate-normalized data from the cat retina (gray + black spikes) did not exhibit dramatic visual differences as a function of stimulus size. However, inspection of the corresponding single-trial frequency spectra, computed using the same set of representative multi-unit spike train segments (black spikes only), once more exhibited distinct peaks in the upper gamma band as a function of increasing stimulus dimensions (Fig. 3c). The largest stimulus, a $9.8° \times 9.8°$ square spot, evoked clear peaks in the upper gamma band, whereas no such peaks were evident in response to the smallest stimulus, corresponding to a $0.7° \times 0.7°$ square spot. As in the retinal circuit model, the increased baseline variability in the spectral amplitudes measured in response to the smallest stimulus may be attributed to the reduced number of spikes in the multi-unit record. Consistent with this interpretation, the variability increased markedly for spike train segments recorded later in the response period (e.g. 0.2 vs. 2.0 s). Likewise consistent

with the results computed from the artificially generated spike trains, the trial-averaged frequency spectra computed from the cat sequences exhibited size-dependent peaks in the upper gamma band (Fig. 4c). For the cat data, the size-dependent peaks computed from the rate-normalized spike trains (gray traces) were smaller than those computed directly from the original recordings (black traces). Because the average firing rate of the cat ganglion cells declined throughout the stimulus, rate-normalized spike trains recorded later in the response required the addition of a proportionately greater number of Poisson distributed events, thereby diluting the peak gamma activity relative to the DC amplitude. As in the retinal circuit model, the rate-normalization process produced a downward shift in the trial-averaged spectral amplitudes evoked by the small spot, which again may be attributed to the reduction in single-trial variability due to the addition of Poisson distributed events. The rate normalization procedure also reduced the amplitude of periodic oscillations in the single-trial correlation functions (Fig. 5c), which although smaller than those computed from the artificial spike trains nonetheless exhibited a clear increase in gamma activity as a function of increasing stimulus size.

The trial-averaged correlations functions computed from the cat data (Fig. 6c), particularly those based on the original sequences (not rate-normalized, black traces), were close in magnitude to the corresponding quantities computed from the model-generated spike trains. Likewise, over the range of spot dimensions examined here, oscillatory modulations increased smoothly with stimulus size for all three data sets and the duration and frequency of the underlying coherent oscillations were similar. The quantitative fit between the model and experimentally derived correlation functions was particularly strong with respect to the amplitudes of the main side peaks. The amplitudes of the central correlation peaks computed from the experimentally recorded spike trains were confounded by the inclusion of the single-cell auto-correlations functions, which were removed explicitly from the model generated correlation functions (see Methods). Another characteristic common to all three data sets was that periodic modulations often persisted longer in the single-trial traces then in the trial-averaged correlations functions, due to the fact that the instantaneous phase of the evoked oscillations drifted randomly over time and thus tended to wash out when computing the mean response over trials. Overall, representative examples of multi-unit data recorded from the cat retina provide additional support for the hypothesis that coherent oscillations can reliably encode global stimulus properties, such as size, on behaviorally meaningful time scales.

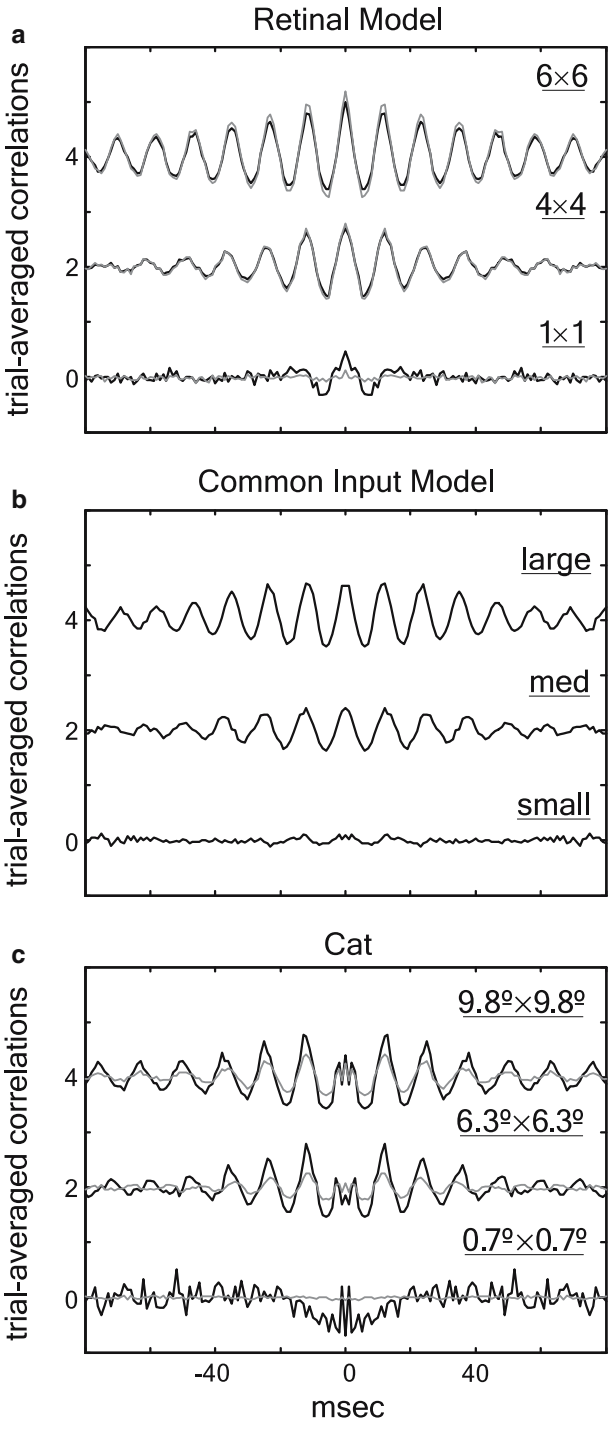

Fig. 6 Trial-averaged correlation functions for different sized spots. Same general organization as in Fig. 3. **a–c** *Black traces* were computed using the original spike trains. *Gray traces* were computed using the rate-normalized data. The amplitude, frequency, persistence and size-dependence of the trial-averaged correlations were similar for all three data sets. Trial-averaged correlations were generally weaker than representative single-trial examples, due to the lack of strong alignment between corresponding maxima. Except for a general dilution of gamma activity in the experimentally recorded sequences, rate-normalization did not qualitatively affect the overall shape of the trial-averaged correlations

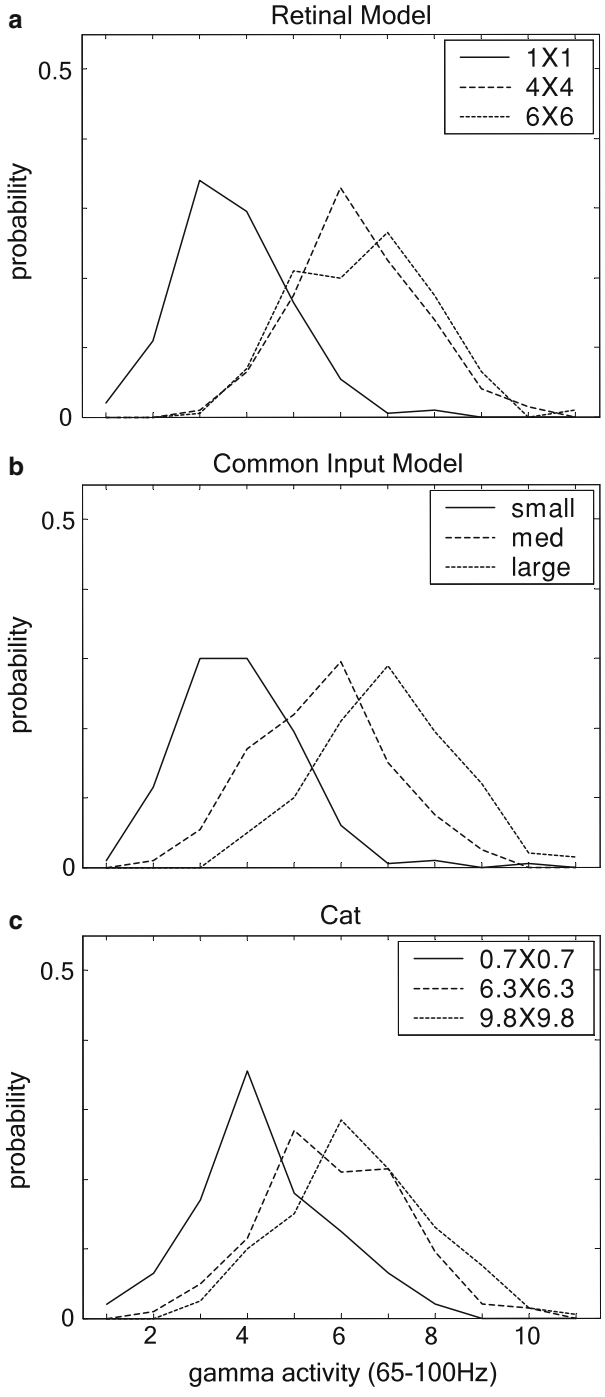

Fig. 7 Distribution of single-trial gamma activity as a function of stimulus size. Gamma activity, defined as the mean Fourier amplitude between 65 and 100 Hz relative to either the DC band (*panel b*) or to the mean amplitude between 220 and 500 Hz (*panels a and c*), was sorted into uniform bins (arbitrary units). **a** Distribution of single-trial gamma activity recorded from the retinal circuit model. **b** Distribution of gamma activity for the common input model. **c** Distribution of single-trial gamma-band activity recorded from cat retinal ganglion cells. **a–c** The distribution of gamma activity elicited by a small spot (*solid line*) was easily distinguishable from that elicited by the intermediate (*dashed line*) and large spots (*dotted line*)

3.2 Size discrimination

To more fully examine the size information encoded by retinal oscillations on physiologically relevant spatiotemporal scales, we constructed probability distributions of the gamma activity present on single-stimulus trials (Fig. 7). The gamma activity was defined as the average spectral amplitude between specified limits, here chosen as 65–100 Hz. Gamma activity was scaled either by the DC component (number of spikes) present on each trial, or when using the original, non-rate-normalized spikes trains, by the average spectral amplitude between 220–500 Hz present during the same 200 ms epoch, as the latter procedure more effectively eliminated the size information encoded by non-linear baseline shifts present at very low firing rates. For each data set and stimulus condition, a measure of gamma activity was computed for all 200 trials and the results sorted into 11 uniform bins spanning the full range of the distribution. We show here only the gamma activity distributions for the original spike trains, although similar results were obtained from the rate-normalized data as well. In all three sets of data, the distributions for the smallest (solid lines) and largest (dotted lines) spot sizes were reasonably distinct, implying that gamma activity from a small group of neighboring ganglion cells can mediate good single-trial size discrimination. On the other hand, the retinal circuit model produced largely overlapping distributions of gamma activity for the intermediate and large spots, an effect that was present in the cat data as well. Although the amplitudes of the spectral peaks evoked by the large spot were greater than those evoked by the intermediate spots, the corresponding spectral widths became narrower, leaving the total area relatively unchanged. We used the average spectral amplitude in the upper gamma band, broadly defined, to assess gamma activity because this quantity seemed more likely to be accessible to downstream neurons than the amplitude of the peak itself. The distributions of gamma activity for the three different spot sizes were better separated in the common input model, probably because the underlying mechanism was more linear. For the cat spike trains, the distribution of gamma activity evoked by the smallest spot size was shifted somewhat to the right compared to the corresponding distributions from the two models, reflecting the higher level of noise present in the single-trial spectra computed from the experimental data. Discrimination between small and large spots is thus expected to be slightly worse for the experimental as compared to that mediated by the model spike trains.

To further quantify the degree of discrimination theoretically possible on single stimulus trials, we employed

a Bayes discriminator (Duda et al. 2001). For any pair of spot sizes, the Bayes discriminator is given by the value of the gamma activity at which the two distributions cross, although this definition does not apply if the distributions have multiple intercepts. If the single-trial gamma activity falls to the left of the Bayes discriminator, the stimulus is classified as the smaller of the two spot sizes, whereas the stimulus is classified as the larger of the two if the gamma activity falls to the right of the discriminator. The overlap between the two distributions estimates the number of trials that will be classified incorrectly. Using the probability distributions described above based on the original (non-rate-normalized) data, we found that for the retinal circuit model, based on 200 m s segments containing spike trains from four neighboring ganglion cells, it was theoretically possible to discriminate between the larger and smaller spot on approximately 85% of the trials (Fig. 8a, black bars). Thus, the gamma activity from a small number of neighboring ganglion cells conveyed enough information on behaviorally relevant time scales to reliably encode a global property of the stimulus, namely size. The common input model mediated qualitatively similar performance on the size discrimination task, albeit with a steeper dependence on the relative dimensions of the two stimuli, again suggesting that the global topological information encoded by coherent oscillations may be largely independent of their mode of generation (Fig. 8a, dark gray bars).

Performance on the size discrimination task mediated by cat retinal ganglion cells was similar to that predicted by the artificially generated spike trains, although somewhat lower, ranging between 70 and 80% (Fig. 8a, light gray bars). The slightly poorer performance mediated by the experimentally recorded data was probably due to the high variability of the gamma activity evoked by the small spot, a direct consequence of the very small number of spikes in each sequence, especially those recorded late in the response period. In general, our quantitative estimates of the size discrimination mediated by the cat data must be carefully assessed with regards to various confounds implicit in the experimental procedure. The experimentally recorded spike trains may have included different ganglion cell types, some of which may not have contributed to the coherent oscillations, and the different sized stimuli may have activated separate, partially overlapping cell populations. The changing baseline of the experimentally recorded spikes trains may also have confounded the discrimination analysis, despite efforts to correct for the effects of adaptation. Thus, the performance levels obtained with the experimentally recorded data, while consistent with our theoretical results, do not represent a quantitative experimental test of the predictions of the retinal model.

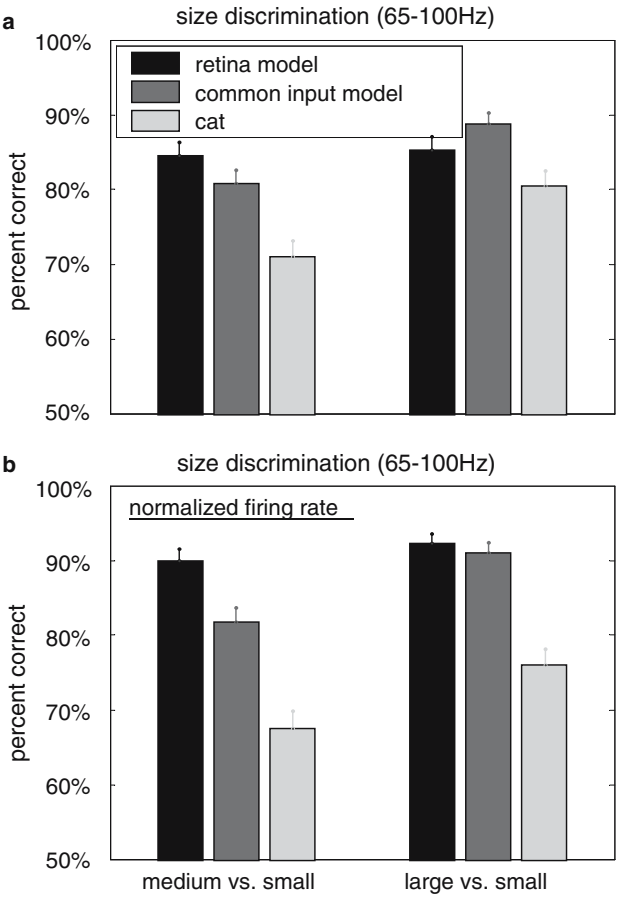

Fig. 8 Size discrimination based on single-trial gamma activity (65–100 Hz). A Bayes discriminator was used to categorize relative spot size, as either "smaller" or "larger", based on the single-trial gamma activity extracted from multi-unit spike trains 200 m s in duration. Two pairs of spot sizes were compared, either small versus medium or small vs. large. The sizes of the small, medium and large spots were as indicated in previous figures. The height of each bar gives the maximum average percentage of trials that could be classified correctly, assuming both sizes were equally likely a priori, based on the distributions shown in Fig. 7. **a** Size discrimination using original spikes trains. **b** Size discrimination using rate-normalized spike trains. **a**, **b** Coherent oscillations within a small group of ganglion cells are capable of supporting reasonable discrimination between small and large spots on behaviorally relevant time scales. Rate-normalization slightly improved size-discrimination in the retinal circuit model by reducing the variability of responses to the small spot, an effect opposed in the cat data by the dilution of gamma activity

Rate-normalized spike trains mediated similar performance on the size discrimination task as was obtained from the original data (Fig. 7b). By reducing the single-trial variability in the responses to the small spot, the rate normalization procedure produced a modest improvement in the size discrimination mediated

by the retinal circuit model. Because rate-normalization was not performed on data from the common input model, slight differences between corresponding bar graphs reflect intrinsic variations in predicted performance arising from the use of different random number sequences and are consistent with the magnitude of the estimated error bars. Unlike the retinal circuit model, rate-normalization did not seem to affect the size discrimination mediated by the experimental data, possibly because any improvements in performance resulting from the lower variability of the responses to the small spot were cancelled by the dilution of the gamma activity evoked by the larger spots. The overall similarity in performance mediated by the original versus the rate-normalized data indicates that reliable size discrimination was possible even when the average number of spikes was held fixed across stimulus conditions.

The range of frequencies included in our definition of gamma activity substantially influenced the ability to discriminate different sized spots. For example, by narrowing the operational range to lie between 70 and 90 Hz, thus more precisely demarcating the size-dependent peaks evident in the frequency spectra, performance on the size discrimination task improved for all three sets of data, exceeding 95% for model generated spike trains and 85% for experimentally recorded sequences (Fig. 9). Again, the rate normalization procedure slightly improved the size discrimination mediated by the retinal circuit model but had little qualitative effect on task performance supported by the cat data. Although even more precise definitions of gamma activity could improve size discrimination still further (results not shown), it becomes correspondingly less clear how downstream neurons could restrict their sensitivity to increasingly narrow frequency bands. Since the issue of how coherent oscillations might be decoded by downstream neurons is outside the scope of the present study, the following analysis employed a fairly conservative definition of gamma activity, using an analysis window of at least 35 Hz.

One motivation for constructing rate-normalized spike trains was to verify, for all three data sets, that coherent oscillations could reliably encode global stimulus properties even when the number of spikes conveyed no information about stimulus size. That the firing rate indeed conveyed no global topological information was shown explicitly by repeating the size discrimination task using the total spike count, instead of gamma activity, as the stimulus classification measure (Fig. 10a). Using rate-normalized spike trains, the total spike count on single stimulus trials supported performance levels that were no better than chance to within statistical fluctuations.

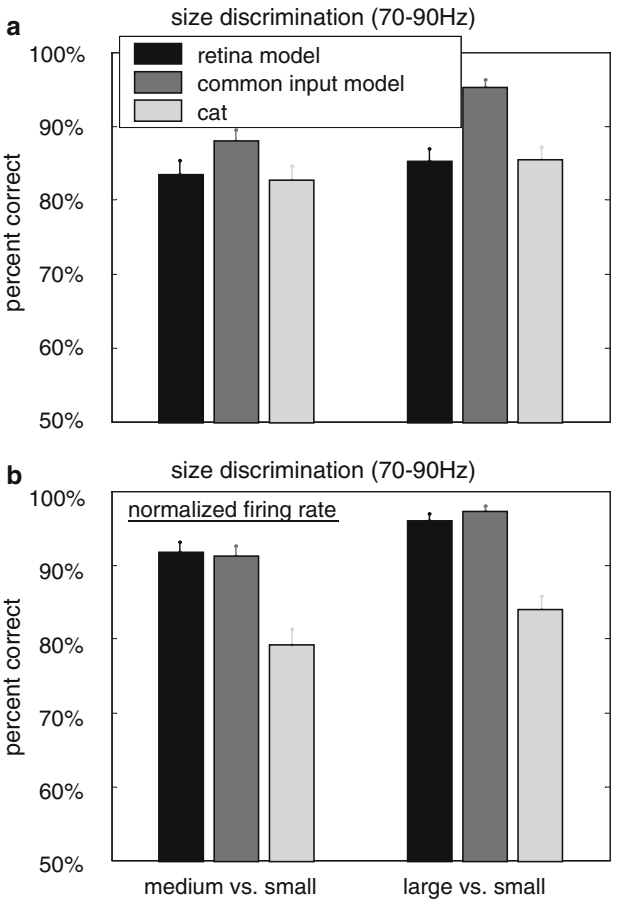

Fig. 9 Size discrimination based on single-trial gamma activity (70–90 Hz). Same organization as in Fig. 8. **a**, **b** Narrowing the definition of gamma activity to more tightly encompass the size-dependent peaks in the frequency spectra yielded improved performance on the size discrimination task

The second motivation for constructing rate-normalized spike trains was to compare the global topological information conveyed by coincident events with that encoded by coherent oscillations. This comparison could be conducted directly by assessing performance on the same size discrimination task using an identical set of rate-normalized spike trains. We found that the total number of coincident events mediated relatively poor performance on the size discrimination task, the maximum percentage of correctly classified trials being approximately 70% (Fig. 10b). For this analysis, a coincident event was defined as two or more spikes arriving in the same 1 m s time bin. We explored other bin widths and/or thresholds but failed to obtain improved performance (data not shown). Because some synchronous events could not be resolved in the experimental data, the performance estimates given here for the cat spike trains should be interpreted cautiously. Nonetheless, the dramatic difference in size discrimination mediated by

coherent oscillations as compared to that mediated by the total number of coincident events, a difference that persisted across both model and experimentally recorded data sets, suggests that gamma activity may be a more reliable indicator of global stimulus properties, at least as measured locally from a few neighboring cells, than is temporal synchrony. The finding that temporal synchrony does not provide a reliable size-dependent signal is moreover consistent with a casual inspection of the representative raw spike trains shown in Fig. 2, which indicate that coincident events are present in the responses to small spots across all three data sets. Whereas the relatively high baseline rate of chance coincident events leads to impaired size discrimination mediated by temporal synchrony, the absence of coherent oscillations in the responses to small spots leads to improved size discrimination based on gamma activity.

We next examined how the information encoded by the coherent oscillations between neighboring retinal neurons depends both on the spatial and temporal extent of the pooled response. Performance on the size discrimination task was again measured by the percentage of correctly classified trials, such that the small and large spots were presented with equal probability. Gamma activity was defined as the average spectral amplitude between 60 and 95 Hz, scaled as described previously. Only the original, non-rate-normalized spike trains were used for this portion of the study. Performance was plotted as a function of the duration of the multi-unit spike train segments incorporating different numbers of model ganglion cells, arranged in groups of either $1 \times 1, 2 \times 2, 3 \times 3$, or 4×4 (Fig. 11a). Performance on the size discrimination task ranged from a low of between 55 and 65% for spike trains consisting of only a single model ganglion cell, to highs of approximately 95% or better for spike trains consisting of between 9 and 16 model ganglion cells and lasting from 200 to 400 m s in duration. Even for analysis windows lasting only 100 m s, performance levels of 75 and 85% were obtainable with 9 and 16 cells, respectively. For the shortest analysis window considered here, equal to 50 m s, maximum performance using spike trains from the retinal model fell to 75%. For all window sizes, the degree of discrimination increased systematically as more neurons were included in the multi-unit spike train. For a fixed number of neurons, performance generally increased as a function of the duration of the spike train segments as well, although there were a few exceptions to this rule, possibly resulting from ongoing random drift in the phase of the coherent oscillations, making the extraction of gamma activity more difficult over long intervals.

Artificial spike trains generated by the common input model, consisting of 1, 4, 9, or 16 units, also mediated

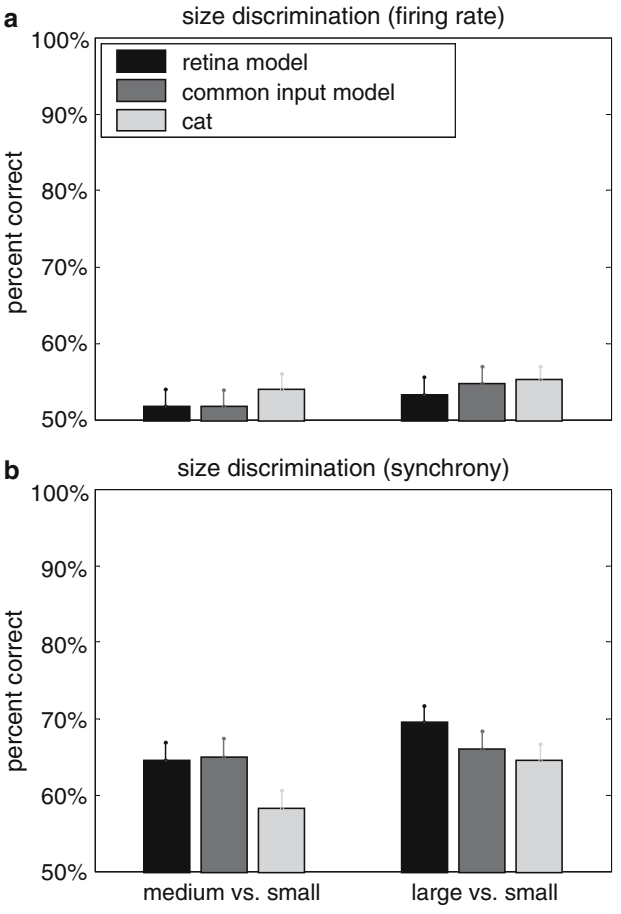

Fig. 10 Size discrimination based on coincident events. **a** The size discrimination task was repeated using the total number of spikes in the rate-normalized event trains to classify stimuli as either smaller or larger. As expected, performance was near chance. **b** Size discrimination using the total number of coincidences (two or more events in a single 1 m s time bin). Performance was better than chance but substantially worse than that mediated by coherent oscillations

improved performance on the size discrimination task as both the duration and number of units increased. The common input model mediated comparatively better discrimination over very short intervals, 50 m s in duration, a discrepancy that may be related to the synchronization of neighboring ganglion cells in the retinal network model due to extensive gap junction coupling via amacrine cells (Kenyon et al. 2003). Nonetheless, the results from both models suggest that the information conveyed by coherent oscillations, especially on physiologically relevant spatial and temporal scales, depends strongly on both the duration of, and number of cells included in, the multi-unit spike train.

We have so far focused on relatively coarse discriminations between large and small spots. Model spike train were used to study whether coherent oscillations could,

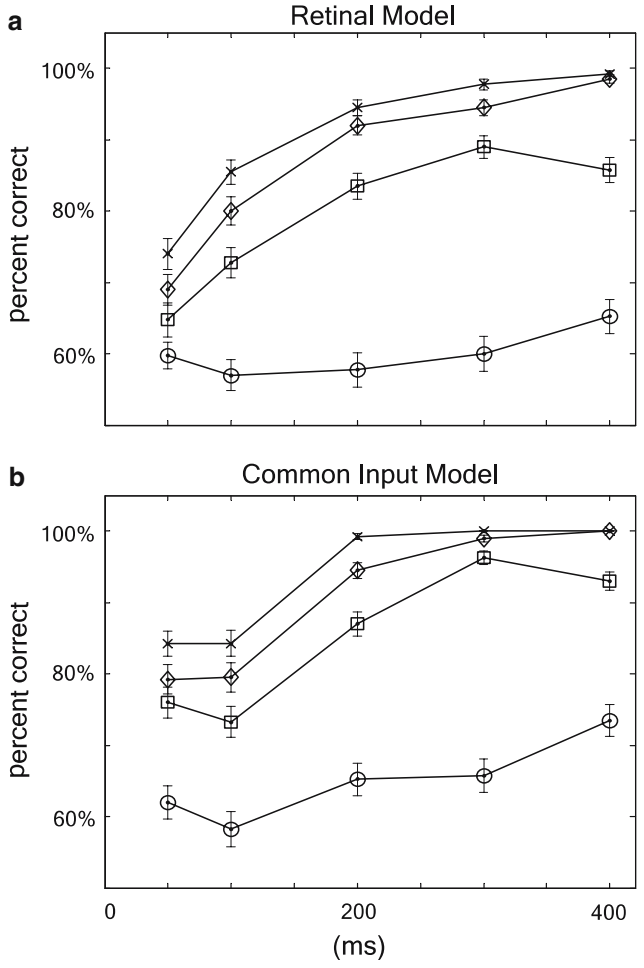

Fig. 11 Size discrimination improves as multi-unit spike train segments become longer and/or include more cells. Analysis was performed using original (non-rate-normalized) event trains with gamma activity defined as the average spectral amplitude between 60 and 95 Hz, scaled as described previously. A Bayes discriminator was used to classify spots as being either "smaller" or "larger" based on the single-trial gamma activity. **a** Percentage of correctly classified trials plotted as a function of the analysis window size for multi-unit spike trains recorded from the retinal model. (Explanation of symbols: *open circle*=1 × 1, *open square*=2 × 2, *open diamond*=3 × 3, × =4 × 4 model ganglion cells). **b** Percent of correctly classified trials using multi-unit spike trains generated by a rate-modulated Poisson process. (Explanation of symbols: *open circle* = 1, *open square* = 4, *open diamond* = 9, × = 16 units). **a**, **b** Performance improved as the analysis window size and the number of neighboring ganglion cells included in the analysis was increased

in principle, support finer discriminations over a wider range of spot sizes, from 1 × 1 to 16 × 16 ganglion cells (Fig. 12). The size of the smaller spot in each pair is given by the intercept of each curve with the *x*-axis, whereas the size of the larger spot is given by the *x*-coordinate of each point. Moving left to right along each curve, the size of the smaller spot is held constant while the size of the larger spot increases as indicated. The analysis window for all pairwise size discriminations was 200 m s. Gamma activity was defined as the mean spectral amplitude between 60 and 95 Hz and only the original, non-rate-normalized spike trains were used for this analysis. Our results suggest that the size information encoded by coherent oscillations, at least on behaviorally relevant time scales, is somewhat binary, allowing discrimination only between small and large spots. This coarse encoding was evident whether the multi-unit spike trains used in the discrimination analysis contained 4 cells (left column) or 9 cells (right column). Only poor discrimination (< 75%) was possible between pairs of spots whose diameters both exceeded approximately 4–6 ganglion cells. While these results await detailed experimental confirmation, they suggest that the size information conveyed by high frequency coherent oscillations among retinal neurons is something of an either-or proposition, at least as measured by the average gamma activity, permitting discrimination only between small and large spots. It is possible that additional size information might be encoded by retinal oscillations if defined more narrowly, such as by the peak amplitude, rather than the average power, in the upper gamma-band, although it is less clear how such information could be extracted by downstream neurons.

3.3 Shape independence

Finally, we addressed the question of whether size-dependent retinal oscillations, for a fixed level of contrast, are independent of object shape. The answer to this question may have important behavioral implications. For example, could retinal oscillations be used, at an early processing level, to distinguish a swarm of small, disconnected fragments (i.e. flies) from a large contiguous object (i.e. a bird)? To investigate this issue, we examined the responses of the retinal model to sets of random binary images. If retinal oscillations are indeed shape independent and sensitive only to contiguous size, then the network responses to a set of random binary images should exhibit a sharp transition in total gamma activity as the average pixel density exceeds a critical percolation threshold at which large connected clusters first appear, due to a percolation phase transition in the 2D image space (Grimmett 1999). In the present context, a connected cluster is defined as a group of ON pixels all connected by at least one path, regardless of length, that includes no OFF pixels. Paths are defined in the site sense of nearest neighbor. There are slight differences between site and bond percolation (e.g. different critical exponents) that are overwhelmed by lateral smoothing in the retina and don't concern us here.

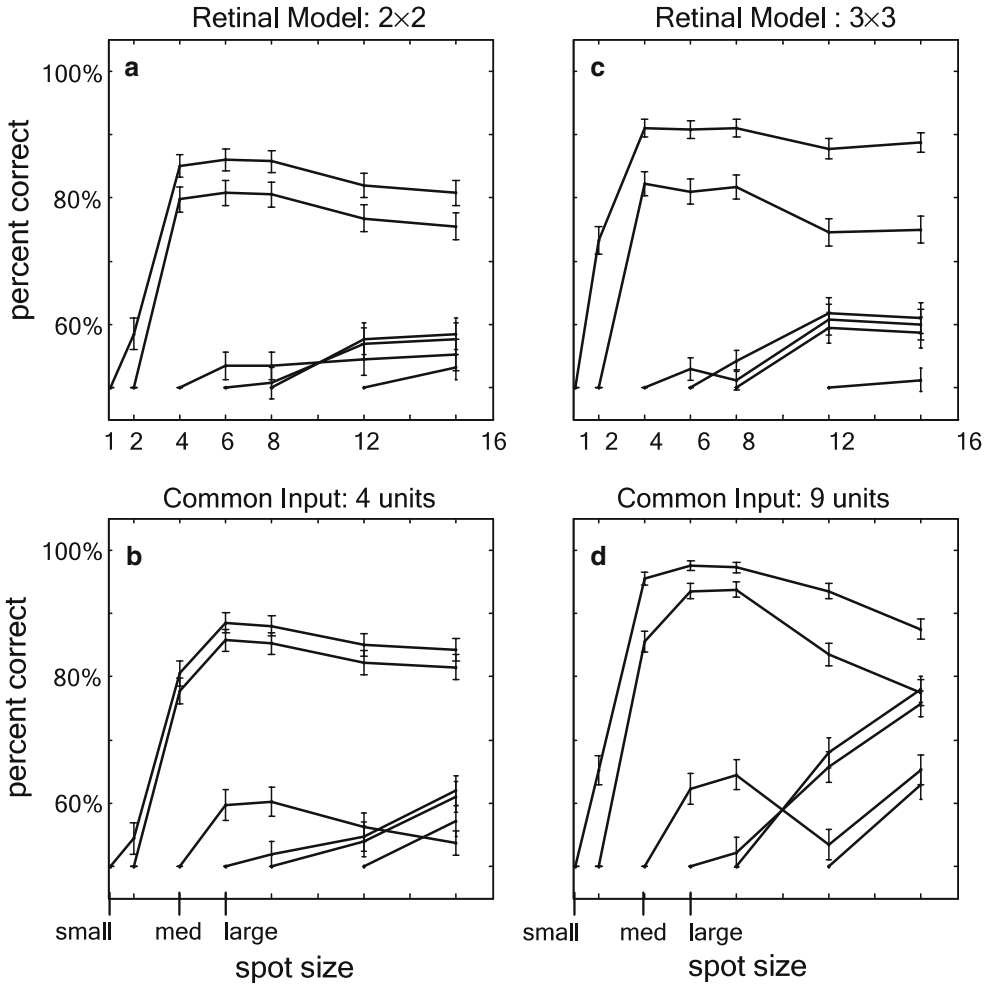

Fig. 12 Coarse discrimination between spot sizes. Pairs of stimuli, consisting of two different sized spots, were classified as either "smaller" or "larger" based on the single-trial gamma activity, falling within the range 60–95 Hz, computed from non-rate-normalized data. The maximum percent correct, estimated via Bayesian discrimination, is plotted for all pairwise combinations. Each point represents a different discrimination task. The size of the smaller spot is given by the x-intercept of the continuous curve and the size of the larger spot by the x-coordinate. **a**, **c** Multi-unit spike trains from the retinal model consisting of **a** 2×2 and **c** 3×3 neighboring ganglion cells. **b**, **d** Rate-modulated Poisson process (common input model) consisting of **b** 4 units and **d** 9 units. Spot sizes specified as L-values (see Methods) ranging from 1 to 16. Sizes corresponding to the labels "small", "med" and "large" used in previous figures are indicated separately. **a**–**d** Performance on the pairwise size discrimination task was poor between spots containing more than 4–9 cells, suggesting that coherent oscillations primarily allow only coarse discriminations between small and large spots

A set of random binary images was constructed by independently turning each pixel white, against a black background, with a probability p between zero and one (Fig. 13a). We computed trial-averaged frequency spectra from massed spike trains that included all 128×128 ganglion cells in the model (the size of the model was increased for these experiments to reduce finite-size effects). At low pixel densities, the image consists entirely of small isolated clusters (Fig. 13a, $p = 0.1$). Because the retinal oscillations produced by separate objects are not coherent (Ishikane et al. 1999; Kenyon et al. 2004c; Neuenschwander and Singer 1996), the total gamma-band activity remains negligible at low pixel densities, at which the oscillations produced by the separate clusters possess random relative phases and thus tend to cancel out (Fig. 13b, $p < p_c$ with $p_c \approx 0.6$ for site percolation). On the other hand, for pixel densities at or above the percolation threshold there exist connected clusters that span the entire image and across which the oscillations are coherent. As expected, there was a sharp increase in mean spectral amplitude in the upper gamma-band as the pixel density crossed the percolation threshold, implying that coherent retinal oscillations encode the presence of large features regardless of their precise shape (Fig. 13b, $p > p_c$). Both the average gamma-band activity (Fig. 13c, solid line) and the average cluster size

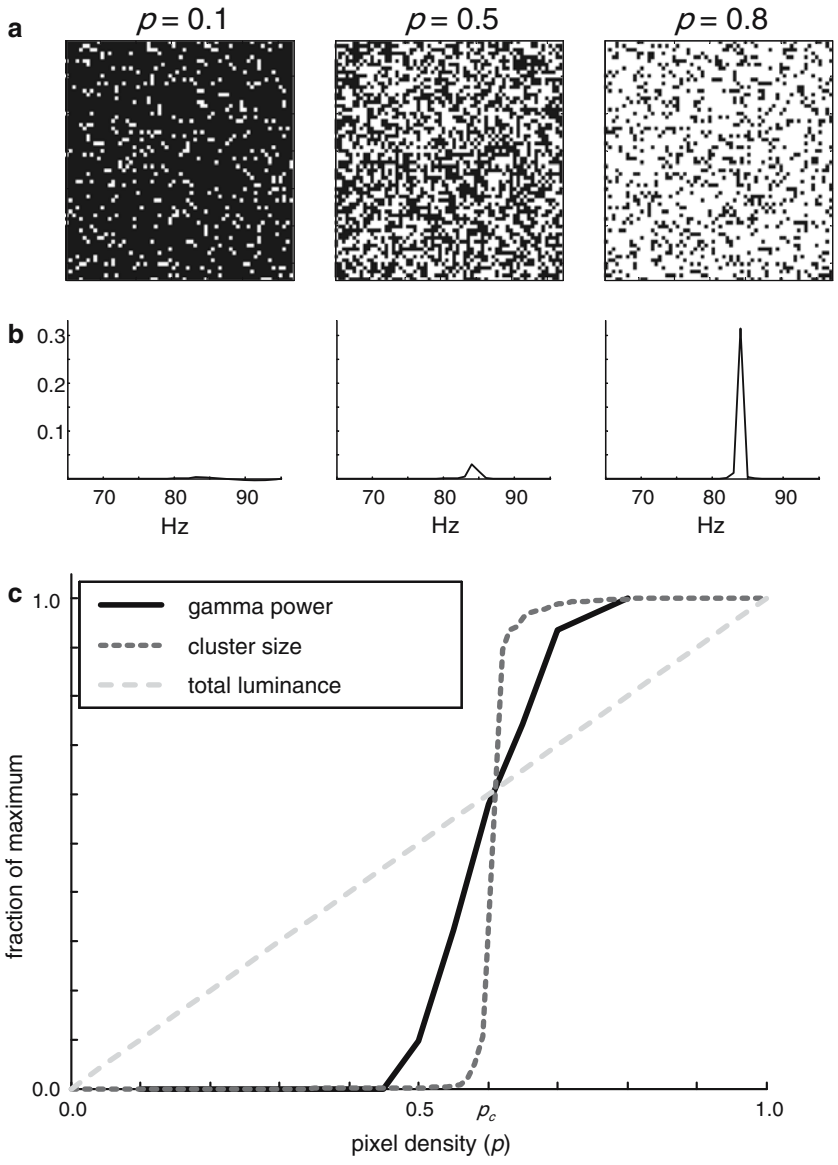

Fig. 13 Binary images produce a sharp transition in gamma activity at the percolation threshold. **a** Random binary images with varying pixel density, p, indicated above each image. Above the percolation threshold ($p \approx 0.6$), connected clusters of white pixels span the image. For clarity, the total number of pixels was reduced in these illustrations from 256×256 to 64×64. **b** Trial-averaged multi-unit frequency spectra computed from massed spike trains combining all 128×128 model ganglion cells. Above the percolation threshold, a large peak appears in the upper gamma band. **c** Gamma activity (*solid black line*) exhibits a sharp transition around the percolation threshold, consistent with the phase transition in the mean size of connected clusters (*dotted dark gray line*). Total luminance (*dashed light gray line*), which grows linearly with pixel density, cannot account for the sharp transition in gamma activity. There were slight differences in the implementation of the retinal circuit model used for these experiments (Miller et al. 2006), but these had no qualitative effect on our conclusions

(Fig. 13c, dotted dark gray line, Hoshen and Kopelman 1976) exhibited sharp transitions around the percolation threshold, whereas the average luminance was simply proportional to the mean pixel density (Fig. 13c, dashed light gray line). Retinal processing is thus predicted to convert a phase transition in the spatial properties of the input image into a transition in the temporal patterning of the output spike trains.

The fact that the mean luminance was not held constant as a function of pixel density does not appear to seriously confound our results. A sharp increase in gamma activity was still evident near the percolation threshold even when the mean luminance was kept constant by reducing the intensity of each pixel in proportion to the mean pixel density (not shown), although the transition in this case was smoother due to the strong dependence of retinal oscillations on stimulus intensity or contrast (Kenyon et al. 2004b; Neuenschwander et al. 1999). The network model lacked known physiological mechanisms for luminance adaptation that would have likely reduced the importance of mean image intensity. The transition in gamma activity was not as sharp as the transition in mean cluster size, presumably because lateral processing in the retina tends to lump together large, nearly connected clusters (Kenyon et al. 2004c). Our retinal model contained only ON-ganglion cells, but we predict similar phenomenon to be exhibited by OFF-ganglion cells in response to contrast-inverted stimuli.

4 Discussion

Using artificial spike trains generated by a computer model of the inner retina, we examined the role of coher

ent oscillations in the encoding of visual information. Our findings support the hypothesis that coherent high frequency oscillations allow small clusters of neighboring ganglion cells to convey global stimulus properties, such as size, over biologically reasonable time scales. The ability to represent information in a small number of cells over a single inter-saccade interval is essential if coherent oscillations between retinal neurons are to convey behaviorally relevant stimulus attributes to early visual processing stages where spatial convergence is generally low (Alonso et al. 2001; Usrey et al. 1999).

A mechanism for locally encoding the presence of large visual features early in the visual system could have behavioral advantages. For many organisms, large objects represent possible danger and early detection allows for an early evasive response. It is possible, for instance, that high frequency oscillations between dimming detectors in the frog retina contribute to just such an early warning system (Ishikane et al. 2005). Similar multiplexing of local and global stimulus properties might be present in other sensory systems as well. For example, coherent oscillations allow discrimination between prey and communication signals in weakly electric fish (Doiron et al. 2003) and in the olfactory system encode contextual category information (Friedrich et al. 2004).

The question of how the information encoded by coherent oscillations among retinal ganglion cells might be decoded by downstream networks is beyond the scope of the present study. However, we may imagine that corresponding high frequency resonances, of cellular, synaptic and/or network origin, might allow neural circuits in the LGN or primary visual cortex to respond preferentially to oscillatory input.

Our results show that artificial spike trains generated by two very different mechanisms mediate similar single-trial size discrimination. Coherent oscillations in the retinal model were produced by negative feedback in which ganglion cells, via their electrical coupling to axon-beating amacrine cells, played a crucial dynamical role. Alternatively, realistic spike trains could be produced by modulating the instantaneous event rates of a set of Poisson generators using a common oscillatory input. Whereas the free parameters in both models were chosen to reproduce the trial-averaged correlations measured experimentally, is was not immediately obvious that the single-trial signal-to-noise in the two models would be identical, especially signals conveyed by high frequency oscillations. However, both models supported similar performance on size discrimination tasks, suggesting that the information conveyed by coherent oscillations is not strongly dependent on how such oscillations are produced. While the different methods used here to generate oscillatory spike trains did not span the space of all possible mechanisms, the two models were nonetheless highly complementary. By using both models to test our main hypothesis, we were thus able to increase the generality of our conclusions.

To test the consistency of our main hypothesis against experimental data, we reanalyzed multi-unit spike trains recorded from the cat retina. Compared to the model data, the experimentally recorded spike trains mediated somewhat lower levels of performance on the size discrimination task, with approximately 75% of the stimuli being classified correctly on the basis of the average gamma activity in multi-unit spike trains 200 m s in duration. The recorded population may not have been homogeneous, however, and the contribution of the different cell types may have varied as a function of both the spot size as well as the time from stimulus onset. More extensive experimental studies, in which simultaneously recorded units are isolated and individually identified, will be required before questions regarding the information encoded by high frequency oscillations can be satisfactorily resolved. Nonetheless, our reanalysis was able to confirm that the predictions of the retinal model are at least qualitatively reasonable and therefore potentially worthy of additional experimental study. It may also be possible to test some of the predictions of the retinal model using local field potentials, or even population averages such as the electroretinogram (ERG), as our prediction of a sharp transition in total gamma activity as the pixel density crosses a percolation threshold requires only a general assessment of retinal output at high temporal resolution.

The role of correlations in the retinal code is controversial. Using an information theoretic analysis it has been argued that correlations between retinal ganglion cells encode relatively small amounts of additional information about natural visual scenes compared to their independent firing rates (Nirenberg et al. 2001). However, these results were based on short sections of spike train data, 10 m s in duration, which although sensitive to precise temporal synchrony were too short to resolve oscillations in the upper gamma band. We similarly failed to find strong evidence that synchrony, as measured by the additional number of coincidences among a few neighboring ganglion cells, provided a reliable encoding of global features such as spot size. However, our results do not address other types of information that might be encoded by synchronous events (Schnitzer and Meister 2003). The approach taken here, namely, examining what stimulus properties can be inferred over physiologically meaningful spatial and temporal scales, suggests a reasonable strategy for resolving such questions. In addition, the role of synchrony

as a possible binding mechanism has sparked intense debate (Shadlen and Movshon 1999; Singer and Gray 1995). Our results suggest that coherent oscillations may encode global stimulus information regardless of their possible contribution to feature binding. Finally, it has been observed that any correlation between neural spike-trains is detrimental to rate-coded signals (Shadlen and Newsome 1998). Correlations, however, are ubiquitous throughout the brain and our results demonstrate that, at least in some cases, may convey useful information not unambiguously represented by local time-averaged firing rates.

Acknowledgements The authors wish to thank Chris Wood, James Theiler, David Marshak and Bryan Travis for ongoing discussions and technical support. This LDRD Program at the Los Alamos National Laboratory provided major funding for this work.

Appendix A. Retinal model

Model spike trains with realistic spatiotemporal correlations were generated by a retinal feedback circuit (Fig. 1), organized as a 32×32 array with wrap-around boundary conditions containing five distinct cell types: Bipolar cells, small amacrine cells, large amacrine cells, axon-bearing amacrine cells, and ganglion cells. All cell types were modeled as single compartment, RC circuit elements obeying a first order differential equation that can be written efficiently in terms of matrix multiplications:

$$\dot{\overleftrightarrow{V}}^{(k)} = -\frac{1}{\tau^{(k)}}\left(\overleftrightarrow{V}^{(k)} - \sum_{k'} \overleftrightarrow{W}^{(k,k')} \cdot f^{(k,k')}\left(\overleftrightarrow{V}^{(k')}\right) \times \overleftrightarrow{W}^{(k,k')^{\mathrm{T}}} - b^{(k)} - \overleftrightarrow{L}^{(k)}\right), \quad (2)$$

where $\overleftrightarrow{V}^{(k)}$ is a 2-D array denoting the normalized membrane potentials of all cells of type k, $(1 \le k \le 5)$, $\tau^{(k)}$ is the time constant, $b^{(k)}$ is a bias current for setting the resting potential, $\overleftrightarrow{L}^{(k)}$ is an external input intended to represent light stimulation, $\overleftrightarrow{W}^{(k,k')}$ gives the connection strengths between presynaptic $\{k'\}$ and postsynaptic $\{k\}$ cell types as a function of their separation along one direction, defined here as 'vertical', $\overleftrightarrow{W}^{(k,k')^{\mathrm{T}}}$ gives the same information as a function of separation along the perpendicular direction, defined here as 'horizontal', and the functions $f^{(k,k')}$ give the associated input-output relations for the indicated pre- and post-synaptic cell types, detailed below. All membrane potentials were subject to a lower cutoff, equal to -1.5. The output of the axon-mediated inhibition was delayed by 2 m s, except for the axonal connections onto the axon-bearing amacrine cells, which was delayed for 1 m s. All other synaptic interactions were delayed by 1 m s. All equations were integrated in Matlab® using a direct Euler method with an integration time step of 1 m s.

The input–output function for gap junctions was given by the identity:

$$f^{(k,k')}\left(\overleftrightarrow{V}^{(k')}\right) = \overleftrightarrow{V}^{(k')}, \quad (3)$$

where the dependence on the presynaptic potential has been absorbed into the definition of $\tau^{(k)}$. This is possible because both the decay term in Eq. 2 and the omitted dependence on the presynaptic potential in Eq. 3 depend linearly on $\overleftrightarrow{V}^{(k)}$, allowing the coefficients to be combined. The input–output function for non-spiking synapses was constructed by comparing, on each time step, a random number with a Fermi-function:

$$f^{(k,k')}\left(\overleftrightarrow{V}^{(k')}\right) = \theta\left(\left[\frac{1}{1+\exp\left(-\alpha\overleftrightarrow{V}^{(k')}\right)}\right] - r\right), \quad (4)$$

where α sets the gain (equal to 4 for all non-spiking synapses), r is a uniform random deviate equally likely to take any real value between 0 and 1, and θ is a step function, $\theta(x) = 1,\ x \ge 0; \theta(x) = 0,\ x < 0$.

Lastly, the input–output relation used for spiking synapses was:

$$f^{(k,k')}\left(\overleftrightarrow{V}^{(k)}\right) = \theta\left(\overleftrightarrow{V}^{(k')}\right). \quad (5)$$

A modified integrate-and-fire mechanism was used to model spike generation. A positive pulse (amplitude = 10.0) was delivered to the cell on the time step after the membrane potential crossed threshold, followed by a negative pulse (amplitude = 10.0) on the subsequent time step. This resulted in a 1 m s action potential that also produced impulse responses in electrically coupled cells, an important element of the circuit dynamics. The bias current, b, was incremented by 0.5 following each spike, and then decayed back to the resting value with the time constant of the cell, adding to the relative refractory period. There was in addition an absolute refractory period of 1 m s.

Along both the horizontal and vertical directions, synaptic strengths fell off as Gaussian functions of the distance between the pre- and post-synaptic cells. For a given horizontal separation, the horizontal weight factor was determined by a Gaussian function of the following form:

Table 1 Cellular parameters

	τ	b	$n \times n$	d	σ
BP	10.0	−0.0	64 × 64	0.25	0.25
SA	25.0	−0.5	64 × 64	0.25	0.25
LA	20.0	−0.25	32 × 32	1.0	0.5
PA	5.0	−0.025	64 × 64	0.25/9.0[a]	0.25/3.0[a]
GC	5.0	−0.025	32 × 32	1.0	0.5

Explanation of symbols: τ Time constant (m s), b bias, $n \times n$: array size, d cutoff radius, σ Gaussian radius (see Eq. 6)
[a] Inner radius/outer radius

Table 2 Synaptic weights

	BP	SA	LA	PA	GC
BP	*	−0.375[b]	−3.0b	−3.0[b] /−15.0[c]	*
SA	3.0[b]	*	−3.0[b]	0.0[b] /−15.0[c]	*
LA	3.0[b]	*	0.25[a]	−3.0[a] /−15.0[c]	*
PA	0.75[b]	−0.75[b]	*	0.25[a] /−45.0[c]	0.25[a,d]
GC	9.0[b]	−4.5[b]	−4.5[b]	0.25[a] /−270.0[c]	*

Each term represents the total integrated weight from all synapses arising from the corresponding presynaptic type (columns) to each cell of the corresponding postsynaptic type (rows), (the quantity $W^{(k,k')}$ in Eq. 4). Asterisk (*) indicates absence of corresponding connection
Synapse type indicated by superscript: [a]gap junction, [b]non-spiking synapse, [c]spiking synapse, [d] Maximum coupling efficiency (ratio of post- to pre-synaptic depolarization) for this gap junction synapse: DC = 11.3%, action potential = 2.7%

$$W^{(k,k')}_{i^{(k)},j^{(k')}} = \alpha\sqrt{W^{(k,k')}}\exp\left[-\frac{\left\|i^{(k)}-j^{(k')}\right\|^2}{2\sigma^{(k,k')^2}}\right] \tag{6}$$

where $W^{(k,k')}_{i^{(k)},j^{(k')}}$ is the horizontal weight factor from presynaptic cells of type k' located in the jth column to the postsynaptic cells of type k located in the ith column, α is a normalization factor, determined numerically, which ensured that the total synaptic input integrated over all presynaptic cells of type k' to every postsynaptic cell of type k equaled $W^{(k,k')}$, $\sigma^{(k,k')}$ is the Gaussian radius of the interaction, and the quantity $\left\|i^{(k)}-j^{(k')}\right\|$ denotes the horizontal distance between the pre- and post-synaptic cells, taking into account the wrap around boundary conditions employed to mitigate edge effects. An analogous weight factor describes the dependence on vertical separation. Equation 6 was augmented by a cutoff condition that prevented synaptic interactions beyond a specified distance, determined by the radius of influence of the presynaptic outputs and the postsynaptic inputs, representing the axonal and dendritic fields, respectively. A synaptic connection was only possible if the output radius of the presynaptic cell overlapped the input radius of the postsynaptic cell. Except for axonal connections, the input and output radii were the same for all cell types. For the large amacrine cells and the ganglion cells, the radius of influence extended out to the centers of the nearest neighboring cells of the same type, producing a coverage factor greater than one (Vaney 1990). The radii of the bipolar, small, and axon-bearing amacrine cells (non-axonal connections only) extended only halfway to the nearest cell of the same type, giving a coverage factor of one (Cohen and Sterling 1990). The external input was multiplied by a gain factor of 3. Values for model parameters are listed in Tables 1 and 2.

Appendix B. Common input model

An oscillatory time series of a duration, T, and temporal resolution, Δt, could be constructed by first defining the discrete frequencies, f_k:

$$f_k = \frac{k}{T}, \quad 0 \le k < \frac{T}{\Delta t} \tag{7}$$

in terms of which the discrete Fourier coefficients were defined as follows:

$$C_k = C\exp(2\pi i r_k)\exp\left(\frac{(f_k - f_0)^2}{2\sigma^2}\right) \tag{8}$$

where f_0 is the central oscillation frequency, σ is the width of the spectral peak in the associated power spectrum, r_k is a uniform random deviate between 0 and 1 that randomized the phases of the individual Fourier components (generated by the Matlab® intrinsic function RAND) and C is an overall scale factor. The coefficients, C_k, were used to convert back to the time domain using the discrete inverse Fourier transform:

$$R_n = A\frac{1}{N}\sum_{k=1}^{N-1} C_k \mathrm{e}^{-2\pi i f_k t_n} + R_0 \quad (9)$$

where the real part of R_n denotes the value of the time-dependent firing rate at the discrete times, $t_n = n\,\Delta t, N = T\Delta t, A$ is an empirically determined scale factor and we have added a constant offset, R_0, which sets the mean firing rate, equal to 50 Hz. The quantity A, with units of Hz, was determined by the formula: $A = R_0(0.16 + 0.14L)$, where L represents the linear stimulus size, or length, along one axis, with values ranging from 1 to 16, and the other coefficients were determined empirically to produce a reasonable match to the experimental data. Stimulus sizes denoted "small", "med" and "large" in the text corresponded to L values of 1, 4 and 6, respectively. The quantity C in Eq. 8 was determined by setting the standard deviation of R_n over all time steps to unity when $A = 1$. The width of the frequency spectrum, σ, also in units of Hz, was given by an analogous formula: $\sigma = 9.4–0.6\ L$. Negative values of R_n were truncated at zero and the resulting time series rescaled so that its average value remained equal to R_0.

The time series defined by R_n was used to generate oscillatory spike trains via a pseudo-random process:

$$S_n = \theta(R_n\Delta t - r) \quad (10)$$

where $R_n\Delta t$ is the probability of a spike in the nth time bin, θ is a step function, $\theta(x < 0) = 0, \theta(x \geq 0) = 1$, and r is again a uniform random deviate. In the limit that $R_n\Delta t \ll 1$, the above procedure reduces to a rate-modulated Poisson process. The same time series, R_n, was used to modulate the firing rate of each element contributing to the artificially generated multi-unit spike train, thus producing temporal correlations due to co-varying input.

References

Alonso JM, Usrey WM, Reid RC (2001) Rules of connectivity between geniculate cells and simple cells in cat primary visual cortex. J Neurosci 21:4002–4015

Ariel M, Daw NW, Rader RK (1983) Rhythmicity in rabbit retinal ganglion cell responses. Vis Res 23:1485–1493

Cohen E, Sterling P (1990) Convergence and divergence of cones onto bipolar cells in the central area of cat retina. Philos Trans R Soc Lond Ser B Biol Sci 330:323–328

Dacey DM, Brace S (1992) A coupled network for parasol but not midget ganglion cells in the primate retina. Vis Neurosci 9:279–290

De Carli F, Narici L, Canovaro P, Carozzo S, Agazzi E, Sannita WG (2001) Stimulus- and frequency-specific oscillatory mass responses to visual stimulation in man. Clin Electroencephalogr 32:145–151

Doiron B, Chacron MJ, Maler L, Longtin A, Bastian J (2003) Inhibitory feedback required for network oscillatory responses to communication but not prey stimuli. Nature 421:539–543

Duda RO, Hart PE, Stork DG (2001) Pattern classification. Wiley, New York

Freed MA (2000) Rate of quantal excitation to a retinal ganglion cell evoked by sensory input. J Neurophysiol 83:2956–2966

Friedrich RW, Habermann CJ, Laurent G (2004) Multiplexing using synchrony in the zebrafish olfactory bulb. Nat Neurosci 7:862–871

Frishman LJ, Saszik S, Harwerth RS, Viswanathan S, Li Y, Smith EL III, Robson JG, Barnes G (2000) Effects of experimental glaucoma in macaques on the multifocal ERG. Multifocal ERG in laser-induced glaucoma. Doc Ophthalmol 100:231–251

Grimmett G (1999) Percolation. Springer, Berlin Heidelberg New York

Hoshen J, Kopelman R (1976) Percolation and cluster distribution. I. cluster multiple labeling technique and critical concentration algorithm. Phys Rev B 14:3438–3445

Ishikane H, Gangi M, Honda S, Tachibana M (2005) Synchronized retinal oscillations encode essential information for escape behavior in frogs. Nat Neurosci 8:1087–1095

Ishikane H, Kawana A, Tachibana M (1999) Short- and long-range synchronous activities in dimming detectors of the frog retina. Vis Neurosci 16:1001–1014

Jacoby R, Stafford D, Kouyama N, Marshak D (1996) Synaptic inputs to ON parasol ganglion cells in the primate retina. J Neurosci 16:8041–8056

Kenyon GT, Harvey NR, Stephens GJ, Theiler J (2004a) Dynamic segmentation of gray-scale images in a computer model of the mammalian retina. In: Tescher AG (ed) Proceedings of SPIE: applications of digital image processing XXVII. Denver 5558:1–12

Kenyon GT, Marshak DW (1998) Gap junctions with amacrine cells provide a feedback pathway for ganglion cells within the retina. Proc R Soc Lond B Biol Sci 265:919–925

Kenyon GT, Moore B, Jeffs J, Denning KS, Stephens GJ, Travis BJ, George JS, Theiler J, Marshak DW (2003) A model of high-frequency oscillatory potentials in retinal ganglion cells. Vis Neurosci 20:465–480

Kenyon GT, Theiler J, George JS, Travis BJ, Marshak DW (2004b) Correlated firing improves stimulus discrimination in a retinal model. Neural Comput 16:2261–2291

Kenyon GT, Travis BJ, Theiler J, George JS, Stephens GJ, Marshak DW (2004c) Stimulus-specific oscillations in a retinal model. IEEE Trans Neural Netw 15:1083–1091

Laufer M, Verzeano M (1967) Periodic activity in the visual system of the cat. Vision Res 7:215–229

Martinez-Conde S, Macknik SL, Hubel DH (2004) The role of fixational eye movements in visual perception. Nat Rev Neurosci 5:229–240

Meister M, Berry MJ II (1999) The neural code of the retina. Neuron 22:435–450

Miller JA, Denning KS, Geroge JS, Marshak DW, Kenyon GT (2006) A high frequency resonance in the responses of retinal ganglion cells to rapidly modulated stimuli: A computer model. Vis Neurosci 23(5) (in press)

Neuenschwander S, Castelo-Branco M, Singer W (1999) Synchronous oscillations in the cat retina. Vis Res 39:2485–2497

Neuenschwander S, Singer W (1996) Long-range synchronization of oscillatory light responses in the cat retina and lateral geniculate nucleus. Nature 379:728–732

Nirenberg S, Carcieri SM, Jacobs AL, Latham PE (2001) Retinal ganglion cells act largely as independent encoders. Nature 411:698–701

O'Brien BJ, Isayama T, Richardson R, Berson DM (2002) Intrinsic physiological properties of cat retinal ganglion cells. J Physiol 538:787–802

Schnitzer MJ, Meister M (2003) Multineuronal firing patterns in the signal from eye to brain. Neuron 37:499–511

Shadlen MN, Movshon JA (1999) Synchrony unbound: a critical evaluation of the temporal binding hypothesis. Neuron 24:67–77, 111–125

Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. J Neurosci 18:3870–3896

Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. Ann Rev Neurosci 18:555–586

Steinberg RH (1966) Oscillatory activity in the optic tract of cat and light adaptation. J Neurophysiol 29:139–156

Stephens GS, Neuenschwander S, George JS, Theiler J, Marshak DW, Singer W, Kenyon GT (2003) See globally, spike locally: retinal oscillations encode large contiguous features. In: Society for Neuroscience Abstracts.

Usrey WM, Reppas JB, Reid RC (1999) Specificity and strength of retinogeniculate connections. J Neurophysiol 82:3527–3540

Vaney DI (1990) The mosaic of amacrine cells in the mammalian retina. In: Osborne NN, Chader GJ (eds)Progress in retinal research. Pergamon Press, Oxford, pp 49–100

Vaney DI (1994) Patterns of neuronal coupling in the retina. Prog Retin Eye Res 13:301–355

Wachtmeister L (1998) Oscillatory potentials in the retina: what do they reveal. Prog Retin Eye Res 17:485–521

Wachtmeister L, Dowling JE (1978) The oscillatory potentials of the mudpuppy retina. Invest Ophthalmol Vis Sci 17:1176–1188