Submitted to: Neural Computation

Correlated Firing Improves Stimulus Discrimination in a Retinal Model

Garrett T. Kenyon[1,4], James Theiler[2,4], John S. George[1,4], Bryan J. Travis[3,4], David W. Marshak[5]

[1]Physics Division (P-21), [2]Non-Proliferation and International Security (NIS-2), [3]Environmental and Earth Sciences (EES-6), [4]Los Alamos National Laboratory, Los Alamos, NM, 87545; [5]Department of Neurobiology and Anatomy, University of Texas Medical School, Houston, TX, 77030

Abbreviated Title: Correlated Firing Improves Discrimination

Corresponding Author:
Garrett T. Kenyon
P-21 Biological and Quantum Physics, MS D454
Los Alamos National Laboratory
Los Alamos, NM 87544
phone: 505-667-1900
fax: 505-665-4507
email: gkenyon@lanl.gov

Acknowledgements: The authors wish to acknowledge useful discussions with Rob Smith and Greg Stephens. This work was supported by the Lab Directed Research and Development Program at the Los Alamos National Laboratory, the Department of Energy Office of Nonproliferation Research and Engineering, and the MIND Institute for Functional Brain Imaging. DWM supported by the National Eye Institute #EY06472 and the National Institute of Neurological Disease and Stroke #NS38310.

Keywords: synchrony, gamma oscillations, rate code, temporal code, retinal ganglion cell, Bayesian discrimination, ideal observer, computer model

Abstract

Synchronous firing limits the amount of information that can be extracted by averaging the firing rates of similarly tuned neurons. Here, we show that the loss of rate-coded information between retinal ganglion cells due to synchronous oscillations can be overcome by exploiting the information encoded by the correlations themselves. Two very different models, one based on axon-mediated inhibitory feedback, the other on oscillatory common input, were used to generate artificial spike trains whose synchronous oscillations were similar to those measured experimentally. Pooled spike trains were summed into a threshold detector whose output was classified using Bayesian discrimination. For a threshold detector with short summation times, oscillatory input yielded superior discrimination of stimulus intensity compared to rate-matched Poisson controls. Even for summation times too long to resolve synchronous inputs, gamma band oscillations in the input spike trains still contributed to improved discrimination by reducing the total spike count variability, or Fano factor. In separate experiments in which neurons were synchronized in a stimulus-dependent manner without attendant oscillations, the Fano factor increased markedly with stimulus intensity, implying that stimulus-dependent oscillations can offset the increased variability due to synchrony alone.

Introduction

Retinal neurons are thought to represent sensory information primarily as changes in their individual firing rates. In order to obtain reliable estimates of rate-coded information on physiological time scales, from 10's to 100's of msec, it may be necessary to estimate the average firing rate over an ensemble of similarly activated cells. Averaging only yields a more accurate estimate of the mean firing rate to the extent that the input spike trains are statistically independent [Mazurek, 2002 #1987; Shadlen, 1994 #1807; Shadlen, 1998 #1802]. When the input spike trains are instead synchronized to some degree, averaging their firing rates no longer produces the same improvement in signal to noise, as fluctuations in the number of spikes arising from individual cells will no longer tend to cancel out across the population of input fibers. Nonetheless, firing correlations are ubiquitous in the nervous system, often associated with coherent oscillations in the gamma frequency band (40-120 Hz) that synchronizes activity both within and between brain areas [Fries, 2002 #1940; Singer, 1995 #1367]. It is thus important to understand what consequences such synchronous oscillations might have for how information is represented by central neurons.

One possibility is that firing correlations simply impose an upper limit on the amount of rate-coded information a population of neurons can represent in its pooled activity over a given unit of time [Mazurek, 2002 #1987; Shadlen, 1994 #1807; Shadlen, 1998 #1802]. From this perspective, firing correlations are an inevitable but undesirable consequence of the massive interconnectivity of neural circuits but otherwise serve no significant information processing function. A second possibility is that the correlations themselves may encode relevant information, a view supported by studies showing that coherent oscillations may be involved in a variety of cognitive processes, including attention [Fries, 2001 #1986], perception [Srinivasan, 1999 #1757], top-down priming [Engel, 2001 #1943], and feature integration [Fries, 2002 #1940; Singer, 1995 #1367]. Theoretical arguments indicate that correlations are not always detrimental to a population code, and can actually increase the overall information content depending on the specific nature of the correlations, tuning differences between the individual neurons, and how the population-coded signals are

extracted [Abbott, 1999 #2260]. Here, we adopt a more empirical approach. Specifically, we employ artificial spike trains exhibiting stimulus-dependent synchronous oscillations similar to those observed between ganglion cells in the cat retina to explore how realistic firing correlations affect the extraction of visual information by a general threshold detection process.

Information transmission was assessed by quantifying the performance of a Bayesian discriminator in classifying different stimulus intensities based on the output of a threshold detector. When driven by correlated inputs in which the stimulus intensity was encoded by both the firing rates of individual cells and by the strength of their synchronous oscillations, threshold elements mediated equal or superior stimulus discrimination than when driven by statistically independent Poisson generators that produced, on average, the same number of spikes. Threshold detectors with short integration times and low background event rates (i.e. coincidence detectors) always extracted more information from spike trains with realistic correlations than from Poisson controls, as the stimulus information encoded by the level of synchrony could be extracted directly. Neurons in the LGN and visual cortex are preferentially sensitive to synchronous inputs [Alonso, 1996 #1370; Usrey, 2000 #1833] and may respond differentially to coincidences in their primary afferents. Somewhat surprisingly, even threshold detectors with long integration times, which were thus unable to resolve synchronous inputs, still mediated equal or superior stimulus discrimination when driven by correlated inputs compared to Poisson controls, consistent with previous theoretical analysis [Abbott, 1999 #2260]. Because the stimulus-evoked oscillations became stronger as a function of the stimulus intensity, the loss of rate-coded information due to the increased firing synchrony was effectively countered by a reduction in the variability of the total input spike count.

Methods

Artificial Spike Trains

Spike trains were generated by one of two computer models, whose parameters were adjusted to qualitatively match the stimulus-dependent, synchronous oscillations recorded from cat ganglion cells [Mastronarde, 1989 #369; Neuenschwander, 1999 #1766; Neuenschwander, 1996 #1331]. Some theoretical studies of correlated activity have employed mathematically generated spike trains in which the degree of firing synchrony was independent of, and thus conveyed no information about, the applied stimulus [Mazurek, 2002 #1987; Shadlen, 1994 #1807; Shadlen, 1998 #1802]. However, the synchronous oscillations between retinal neurons depend strongly on stimulus parameters such as size [Ishikane, 1999 #1934; Neuenschwander, 1999 #1766; Ariel, 1983 #538], contrast [Neuenschwander, 1999 #1766], connectedness [Neuenschwander, 1996 #1331; Ishikane, 1999 #1934] and velocity [Neuenschwander, 1999 #1766]. To produce realistic firing correlations, we adopted two very different yet complementary strategies for generating artificial spike trains. In the first case, we used a detailed model of axon-mediated inhibitory feedback in the inner retina in order to generate synchronous oscillations via a complex, but biologically plausible, dynamical process. In second case, synchronous oscillations were generated using a minimalist approach in which each ganglion cell generated Poisson distributed spikes at a rate modulated by a common source of oscillatory input. Although both models exhibited similar firing correlations, as assessed by comparing their multiunit cross-correlation histograms (multiunit CCHs), defined below, the two models were based on very different underlying mechanisms. Thus, we were able to investigate whether our conclusions regarding the information content of spike trains with realistic correlations were robust with respect to their underlying mode of generation.

To guard against the possibility that our mathematically generated spike trains exaggerated the information conveyed by correlations, both models were subject to the following constraints. First, when corroborating information was available, the largest

stimulus-evoked correlations between the model-generated spike trains were comparable to the correlations measured experimentally under similar conditions. Thus, under these circumstances we are confident that none of the correlations generated by the model were substantially in excess of the correlations present physiologically. Second, the level of correlations in the absence of stimulation was comparable to the spontaneous correlations measured experimentally, ensuring that a correlation code would not possess an unfair advantage by starting from an unphysiologically low baseline level. Third, because the retinal model did not incorporate several known adaptation mechanisms, particularly those occurring in the outer retina, we ensured that the plateau firing rates evoked by stimuli of various intensities were probably larger than those that would occur physiologically. Thus, the model was probably conservative and likely biased in favor of a rate-code.

Retinal Model

Artificial spike trains with realistic firing correlations were generated by a retinal model (fig 1), details of which have been published previously [Kenyon, 2003 #2265]. The model retina was organized as a 32x32 array with wrap-around boundary conditions containing 5 distinct cell types: Bipolar cells (BP), small amacrine cells (SA), large amacrine cells (LA), poly-axonal amacrine cells (PA), and alpha ganglion cells (GC). All cell types were modeled as single compartment, RC circuit elements obeying a first order differential equation of the following form:

$$\dot{V}_{ij}^{(k)} = -\frac{1}{\tau^{(k)}}\left(V_{ij}^{(k)} - \sum_{i'j'k'} W_{ii'}^{(k,k')} \cdot f^{(k,k')}\left(V_{i'j'}^{(k')}\right) \cdot W_{jj'}^{(k,k')} - b^{(k)} - L_{ij}^{(k)} \right), \qquad (1)$$

where $V_{ij}^{(k)}$ is a 2-D array denoting the normalized membrane potentials of all cells of type *k*, $(1 \le k \le 5)$, with *i* denoting the row, and *j* the column, of the corresponding cell, $\tau^{(k)}$ is the time constant, $b^{(k)}$ is a bias current for setting the resting potential, $L_{ij}^{(k)}$ is an external input representing light stimulation, $W_{ii'}^{(k,k')}$ gives the connection strengths between presynaptic, *k′,* and postsynaptic, *k,* cell types as a function of their row (vertical) separation, $W_{jj'}^{(k,k')}$ gives the same information as a function of the column (horizontal) separation, and the functions $f^{(k,k')}$ give the associated input-output relations for the indicated pre- and post-synaptic cell

types, detailed below. The output of the axon-mediated inhibition was delayed by 2 msec, except for the axonal connections onto the axon-bearing amacrine cells, which was delayed for 1 msec. All other synaptic interactions were delayed by one time step, which equaled 1 msec. The ratio of ganglion cells to bipolar cells reflects known density differences. The true ratio is higher, but computational constraints limited the total number of neurons that could be efficiently simulated. However, for the relatively large stimuli used in our experiments, fine details of the ganglion cell receptive field structure should not have been critical. Numerical ratios were always powers of 2, to ensure that the underlying connectivity was translationally invariant between processing modules. All equations were integrated in Matlab using a direct Euler method. Previous analysis has verified that the synchronous oscillations predicted by the model are robust with respect to individual parameter variation and integration step size [Kenyon, 2003 #2265].

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

$$f^{(k,k')}\left(V_{ij}^{(k')}\right) = V_{ij}^{(k')}, \tag{2}$$

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 equation 1 and the omitted dependence on the presynaptic potential in equation 2 depend linearly on $V_{ij}^{(k)}$, allowing the coefficients to be combined.

Many retinal neurons, including bipolar cells and most amacrine cells, do not fire action potentials but rather generate stochastically distributed postsynaptic potentials at a rate proportional to their membrane voltage [Freed, 2000 #1782]. Such stochastic inputs are likely to contribute significantly to ganglion cell membrane noise [van Rossum, 2003 #2264]. The input-output function for graded stochastic synapses was constructed by comparing, on each time step, a random number with a Fermi-function:

$$f^{(k,k')}\left(V_{ij}^{(k')}\right) = \theta\left(\left[\frac{1}{1+\exp\left(-\alpha V_{ij}^{(k')}\right)}\right] - r\right), \tag{3}$$

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, $\theta(x) = 1,\ x \geq 0$; $\theta(x) = 0,\ x < 0$.

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

$$f^{(k,k')}\left(V_{ij}^{(k')}\right) = \theta\left(V_{ij}^{(k')}\right). \tag{4}$$

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 msec action potential followed by an after hyper-polarization that decayed with the time constant of the cell. Action potentials produced impulse responses in electrically coupled cells as well, an important element of the oscillatory feedback dynamics. The bias current, b, was decremented by -0.5 following each spike, and then returned 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 msec.

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 vertical separation, the weight factor was determined by a Gaussian function of the following form:

$$W_{i^{(k)},i'^{(k')}}^{(k,k')} = \alpha\sqrt{W^{(k,k')}}\exp\left[-\frac{\left\|i^{(k)} - i'^{(k')}\right\|^2}{2\sigma^{(k')2}}\right] \tag{5}$$

where $W_{i^{(k)},i^{(k')}}^{(k,k')}$ is the weight factor from presynaptic cells of type k' at column index $i'^{(k')}$ to the postsynaptic cells of type k at column index $i^{(k)}$, where the superscripts k and k' are necessary because the number of rows may be different for the two cell types, α is a numerically determined normalization factor that ensured the total synaptic input integrated over all presynaptic cells of type k' to every postsynaptic cell of type k equaled $W^{(k,k')}$, $\sigma^{(k')}$ is the Gaussian radius of the interaction, which depended only on the presynaptic cell type, and the quantity $\left\|i^{(k)} - i'^{(k')}\right\|$ denotes the vertical distance between the pre- and post-synaptic cells, taking into account the wrap around boundary conditions employed to mitigate edge effects

and assuming all cell types are distributed uniformly over rectilinear grids of equal total area and whose center points coincide. An analogous weight factor describes the dependence on horizontal separation. Equation 5 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, roughly corresponding to 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 then one [Vaney, 1990 #1506]. 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, 1990 #764]. The radius of the axonal connections was equal to 9 ganglion cell diameters. The external input was multiplied by a gain factor of 3, chosen so that a stimulus intensity of 1 would produce an approximately saturating response in the model bipolar cells. Values for model parameters are listed in tables 1 and 2.

Figure 1 about here

Rate-Modulated Poisson Process

A second method of generating artificial spike trains was based on the minimalist assumption that firing correlations between cat ganglion cells could be due to a common oscillatory input. An oscillatory time series of a duration, T, and temporal resolution, Δt, possessing realistic temporal correlations could be constructed by first defining the discrete frequencies, f_k:

$$f_k = \frac{k}{T}, 0 \le k < \frac{T}{\Delta t} \tag{6}$$

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

$$C_k = e^{2\pi i r} \exp\left(\frac{(f_k - f_0)^2}{2\sigma^2} \right) \tag{7}$$

where f_0 is the central oscillation frequency, σ is the width of the spectral peak in the associated power spectrum and r 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). These coefficients were used to convert back to the time domain using the discrete inverse Fourier transform:

$$R_n = A \frac{1}{N} \sum_{k=1}^{N-1} C_k e^{-2\pi i f_k t_n} + R_0 \tag{8}$$

where the real part of R_n denotes the value of the time-dependent firing rate at the discrete times, $t_n = n \cdot \Delta t$, $N = T \cdot \Delta t$, A is an empirically determined scale factor and we have added a constant offset, R_0, which sets the mean firing rate. Values of A reported in the text, given in units of Hz, are expressed relative to the quantity D, the standard deviation of R_n over all time steps with $A = 1$. 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:

$$S_n = \theta(R_n \Delta t - r) \tag{9}$$

where $R_n \cdot \Delta t$ is the probability of a spike in the nth time bin, θ is a step function, $\theta(x<0) = 0$, $\theta(x>0) = 1$, and r is again a uniform random deviate. In the limit that $R_n \cdot \Delta t << 1$, 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 multiunit spike train, thus producing temporal correlations due to shared input.

Correlated Poisson Process

To generate artificial spike trains that possess spatial, but not temporal correlations, we employed a pseudo-Poisson process to generate a template spike train from which mutually correlated spike trains could be constructed. To produce a random template

spike train of duration T, temporal resolution Δt, and mean spike rate R_0, we simply used equation 9 with the replacement $R_n \rightarrow R_0$:

$$S_n^{(0)} = \theta\left(R_0 \Delta t - r\right). \tag{10}$$

To produce spatially correlated spike trains, the template time series, $S_n^{(0)}$, was used to construct new spike trains according to the formula:

$$S_n^{(k)} = \theta\left(S_n \cdot C\Delta t + \left(1 - S_n\right) \cdot \left(1 - C\Delta t\right) \frac{\left(R_0 \Delta t\right)}{1 - R_0 \Delta t} - r \right) \tag{11}$$

where $S_n^{(k)}$ denotes the spike train of the k^{th} cell, C is the conditional firing rate given that a spike occurred in the corresponding time bin of the template train and r is again a uniform random deviate. The maximum value of $C \cdot \Delta t$ was $(1 - R_0 \cdot \Delta t) \cdot 0.5$.

Data Analysis

Reported correlations are expressed as a fraction of the expected synchrony due to chance, measured either during baseline activity or during the plateau portion of the response to a sustained stimulus (200-600 msec). With this normalization, a correlation amplitude of one at zero delay corresponded to a doubling in the number of synchronous events over the expected rate. Correlations were plotted as a function of the delay after averaging over all events occurring during the plateau portion of the response. For each delay value, this average was compensated for edge effects arising from the finite length of the two spike trains. To increase the signal to noise, the firing rates or correlations were averaged over all cells, or distinct cell pairs, responding to the same stimulus, producing a multiunit peri-stimulus-time-histogram (PSTH) or multiunit cross-correlation-histogram (CCH), respectively. Auto-correlation functions were not included in the multiunit CCH, which thus only included cross-correlations between distinct cell pairs. Error bars were estimated by assuming Poisson statistics for the count in each histogram bin. All correlations were obtained by averaging over 200 stimulus trials, using a bin width of 1 msec.

Bayes Discriminator

A Bayes discriminator was used to distinguish between different intensities based on the distribution of threshold events across independent stimulus trials. For each intensity, the number of suprathreshold events was determined on successive trials and the results normalized as a probability distribution. For any given pair of intensities, the percent of correct classifications made by an ideal observer was inversely related to the degree of overlap between the two distributions. Total overlap corresponded to performance at chance (50% correct) while zero overlap implied perfect discrimination (100% correct). For the threshold detection process, all spikes occurring within a given time bin, whose width varied from 1-20 msec depending on the experiment, were summed together and the result compared to a threshold. There was no overlap between successive time bins. The duration of the discrimination interval, which varied from 50 to 400 msec, was adjusted so as to approximately normalize the task difficulty as the number of inputs was varied.

When the discrimination interval contained the transient portion of the response (0-50 msec), constructing an appropriate Poisson control was less straightforward than during the plateau period, during which the firing rate could be treated as constant. Moreover, because high frequency oscillations during the response peak were strongly stimulus-locked, it was not possible to use the multiunit PSTH to estimate the time-dependent event rates of the equivalent Poisson generators, as these would have contained the high-frequency stimulus-locked oscillations that the control is intended to eliminate. By using a boxcar filter with a width of 9 msec, however, we were able to modify the multiunit PSTH so as to remove high-frequency components but without eliminating the central response peak. The average number of spikes was always the same for both the model ganglion cells and the Poisson controls.

Results

The principal characteristics of the retinal model used to generate most of the artificial spike trains used in this study are best illustrated by examining the multiunit PSTHs and multiunit CCHs computed from the simulated responses to a narrow bar stimulus of maximum intensity (fig. 2a). The multiunit PSTH, which combines the responses of all twelve model ganglion cells directly beneath the stimulus, consisted of a sharp peak, about 50 msec wide, followed by a plateau period of sustained elevated activity (fig. 2b). Relative to baseline, the sustained increase in spike activity produced by the bar stimulus was comparable to, or larger than, the sustained increase in firing exhibited by cat ganglion cells in response to high contrast features [Creutzfeldt, 1970 #1777; Enroth-Cugell, 1980 #565]. Likewise, the pronounced downward notch separating the peak and plateau regions is characteristic of ganglion cell responses to large, flashed stimuli [Cox, 1996 #1005].

Despite the absence of periodic structure in the plateau portion of the multiunit PSTH, there were nonetheless very prominent high frequency oscillations in the multiunit CCH recorded during the plateau portion of the response (fig. 2c, solid black line). In the cat retina, high contrast stimuli produce an approximate doubling in the number of synchronous events relative to the expected level due to chance [Neuenschwander, 1999 #1766]. In the spike trains generated by the retinal model, the synchronous oscillations evoked by a maximum intensity bar stimulus resulted in a qualitatively similar increase in the number of synchronous events relative to the expected level. Thus, the artificially generated spike trains used in the present study reflect physiologically reasonable levels of correlated activity.

During the plateau portion of the response, the high frequency oscillations were not strongly phase-locked to the stimulus onset, but rather tended to become phase randomized over time, as revealed by the decline in correlation strength with increasing delay. Consistent with the absence of stimulus-locked oscillations during the plateau period, the shift predictor was negligible (fig. 2c, dashed gray line). Correlations between cat ganglion cells recorded during sustained activity decline in a similar manner with increasing delay and likewise possess negligible shift predictors [Neuenschwander, 1999 #1766].

We also computed the multiunit CCH for the 50 msec period following stimulus onset encompassing the response peak (fig. 2d, solid black line), during which the high frequency oscillations were strongly stimulus-locked, as revealed by the periodic structure in the multiunit PSTH as well as by the large shift predictor (fig. 2d, dashed gray line). A tendency for high frequency oscillations to become less stimulus-locked as a function of time from stimulus onset is also seen in experimental data [Neuenschwander, 1999 #1766].

Figure 2 about here

Firing Correlations and Firing Rate are both Modulated by Stimulus Intensity

To investigate the relationship between correlation strength and stimulus contrast, the same narrow bar covering twelve ganglion cells was varied over a 32-fold range of intensities (fig. 3). Both the peak and plateau firing rates of the model ganglion cells activated by the stimulus, as measured by the multiunit PSTH, increased in a graded manner as the stimulus intensity was raised (fig. 3b). Similarly, the degree of synchrony during both the plateau and peak portions of the response, assessed by the amplitude of the multiunit CCH at zero delay, also increased with stimulus intensity (figs. 3c and 3d, respectively). The amplitude of stimulus-evoked high frequency oscillations between cat ganglion cells depends similarly on luminance contrast [Neuenschwander, 1999 #1766]. As a function of stimulus intensity, synchrony could be modulated over a greater dynamic range, measured relative to baseline activity, than could the fractional change in the multiunit firing rate, both during the response plateau (fig. 3d) and during the response peak (fig. 3e). Overall, firing correlations were very sensitive to stimulus intensity, and thus might convey information about contrast in addition to that represented by the mean firing rate across the ensemble of activated cells.

Figure 3 about here

Correlated Inputs Transmit More Information through Coincidence Detectors than Independent Rate-Matched Controls

A simple threshold detector with a short integration time window and a low rate of suprathreshold background events was able to extract the stimulus intensity more reliably from a hybrid rate/correlation code than from statistically independent, rate-matched Poisson distributed inputs. Spikes from a column of twelve ganglion cells activated by a narrow bar were summed into a simple threshold detector (fig. 4). The event rate of the detector was determined by the total number of times the input crossed threshold during a 200 msec epoch within the plateau portion of the response. For these experiments, the detection threshold was set to a level that required three or more spikes to arrive within the same 2 msec time bin in

order to produce a detector event. To assess the extent to which the detector was able to utilize firing correlations between its inputs, the twelve stimulated ganglion cells were replaced by independent Poisson generators that, on average, produced the same number of spikes per unit time. For both correlated and Poisson input, the baseline detector event rate was very low, around 1 Hz (fig. 4a). A high-intensity stimulus produced a greater increase in the detector event rate when the inputs were correlated as compared to the case when the inputs were independent. This extra sensitivity reflected the fact that a threshold process with a short integration time is well suited for detecting rare synchronous events [Kenyon, 1990 #1375; Kenyon, 1992 #1680], a property also exhibited by cortical and sub-cortical neurons [Usrey, 2000 #1833; Alonso, 1996 #1370].

Figure 4 about here

When driven by correlated input from the model ganglion cells, the output of the threshold detector allowed for better discrimination between different stimulus intensities than when driven by independent Poisson generators. The probability of observing a given number of detection events was plotted for a range of stimulus intensities (fig. 4b). As the stimulus intensity increased, the distributions shifted to the right, reflecting the greater number of suprathreshold inputs. When the threshold detector was driven by correlated input from the retinal model, the event distributions for different intensities were more separable than when driven by the rate-matched controls. To quantify the degree to which firing correlations caused detector output to be more discriminable, we used a Bayes discriminator to estimate the maximum percentage of intensity comparisons that could be correctly classified [Duda, 2001 #1915]. Starting from several different baselines, firing correlations allowed for a higher percentage of correct intensity classifications over a broad range of intensity increments (fig. 4c). The abscissa of each point gives the final stimulus intensity, while the x-intercept of the line passing through that point yields the corresponding baseline intensity. The ordinate of each point gives the percentage of trials on which the given pair of stimulus intensities could be correctly classified, using only the single trial output of the detector. For many pairwise intensity discriminations, firing correlations allowed

approximately ten additional trials out of a hundred to be correctly classified (solid black lines) compared to the Poisson control (dashed gray lines). Averaged over all intensity increments, the extra number of correct classifications was approximately 5/100, but this value is likely conservative due to saturation. Our results suggest that when retinal output is decoded by a threshold detection process with a short integration time window, firing correlations allow information to be transmitted more reliably than would be the case with independent firing rates.

During the peak portion of the response, firing correlations also permitted better discrimination of intensity increments, as represented by the output of the threshold detector over the first 50 msec of the response (fig. 5). However, constructing an appropriate Poisson control was less straightforward than during the plateau period, as the high frequency oscillations occurring during the response peak were strongly stimulus-locked (fig. 5a, left bottom panel, smooth trace). In order to produce control spike trains that contained only firing correlations due to the transient itself, we used a boxcar filter with a width of 9 msec to remove high-frequency oscillations from the multiunit PSTH without eliminating the central response peak. The average number of input spikes arriving during the response peak was the same for both the model ganglion cells and the Poisson control. The distribution of detector events was more separable when input was correlated by high-frequency oscillations than when input was correlated only by the transient response peak (fig. 5b). A Bayesian discriminator analysis again confirmed that on average, 5 additional trials out of a hundred could be correctly classified when input to the threshold detector was correlated by stimulus-dependent, synchronous high frequency oscillations as compared to the non-oscillatory control (fig. 5c). These results show that gamma-band oscillations can convey relevant stimulus information even when superimposed with other sources of firing correlations.

Figure 5 about here

The above results presumably depend on the detailed nature of the firing correlations between the input spike trains. In particular, two sets of input spike trains might possess similar correlations as assessed by the multiunit CCH, which is an average over many independent trials, but still differ significantly with respect to the correlations present on individual trials. Our retinal model provides a biologically plausible method for generating realistic correlations, but as with any model of complex neural circuitry, it is impossible to know to what extent the proposed physiological mechanisms are correct. To control for the possibility that the correlations produced by our retinal model might have inadvertently biased our results, a second set of input spike trains was generated using a very different mechanism. In contrast to the inhibitory feedback dynamics employed in our retinal model, the second model required only a source of common input to modulate, in phase, the firing rates of the stimulated cells. The parameters of the common input model were adjusted so that both approaches yielded similar multiunit CCHs over a range of stimulus intensities (fig. 6a). Despite employing very different underlying mechanisms, both models supported very similar levels of stimulus discrimination, with the difference in total performance being less than one trial in a hundred (fig. 6b). These results indicate that the information conveyed by spatiotemporal correlations is unlikely to be an artifact of a particular mode of generation, but rather reflects a general property of neural populations exhibiting synchronous gamma band oscillations.

Figure 6 about here

Correlations Mediate Greater or Equivalent Stimulus Discrimination for a Wide Class of Threshold Detectors

The hybrid rate/correlation code continued to mediate equal or superior stimulus discrimination even when the integration time of the threshold detector was increased so as to diminish the importance of synchronous inputs. To quantify the performance of the ideal observer for a given threshold detector, we defined the quantity Δ, which gives the difference

in the percentage of correctly classified trials using correlated input compared to the Poisson control, averaged over all intensity pairs. Our results show that correlated input mediates greater or equivalent stimulus discrimination over a wide range of integration times and background activity levels (fig. 7). Plotted as a function of integration time, Δ was largest for small summation intervals well suited to resolve synchronous inputs (fig. 8a). When plotted as a function of the detection threshold, Δ generally increased for integration times less than approximately 10 msec, since larger thresholds produced lower values of background activity and thus made the detection process more sensitive to synchrony (fig. 7b).

A somewhat surprising aspect of our results was that as the integration time became large enough to effectively discard intensity information encoded by the degree of synchrony, Δ did not become strongly negative, but rather approached an asymptotic level near zero. Since synchrony adversely affects the amount of information that can be extracted from the average firing rate over a population of similarly activated neurons [Mazurek, 2002 #1987; Shadlen, 1994 #1807; Shadlen, 1998 #1802], it might have been anticipated that once the integration time became sufficiently long to discount any information encoded by the degree of synchrony, independent Poisson inputs would have mediated greater stimulus discrimination than correlated inputs. However, such reasoning fails to consider the effects of stimulus-evoked oscillations, which like synchrony also increased as a function of stimulus intensity, as indicated by the increased persistence of periodic structure in the multiunit CCH as the stimulus intensity was increased (fig. 3). As a result of the stronger oscillations evoked by higher stimulus intensities, the spike trains became more regular, and thus the total number of inputs over the course of each 200 msec discrimination trail became more reliable predictors of the stimulus intensity.

To quantify the reliability of the afferent spike trains, we computed the Fano factor of the multiunit input as a function of stimulus intensity (fig. 7c). The Fano factor is defined as the variance in the number of spikes divided by the mean, and is equal to one for a Poisson process [Teich, 1989 #1988]. At all intensities, the Fano factor of the combined correlated input (solid line) was less than that of the Poisson control (dashed line), indicating that the

stimulus-evoked oscillations caused the total number of spikes to be less variable, and therefore more reliable, than for independent Poisson generators. In the absence of oscillations, the Fano factor would have increased markedly with stimulus intensity as a consequence of the increased synchrony. To illustrate this fact, a second control was employed in which stimulus-dependent synchrony was introduced in the absence of oscillations. The maximum synchrony was set approximately equal to the maximum synchrony present in the hybrid rate/correlation code, while the minimum synchrony was set to zero and a linear interpolation was used for intermediate intensities. By allowing no more than one spike in each 10 msec time bin, thereby producing a relative refractory period, the Fano factor in the absence of spatial correlations could be reduced to approximately the same level exhibited by the model during background activity. In the absence of oscillations, strong synchrony produced a large rise in the Fano factor of the combined input as the stimulus intensity was increased (dotted line). Our results demonstrate that a refractory period, by itself, cannot account for the reliability of the total integrated input in the presence of strong synchrony, but that such reliability results naturally from the associated high frequency oscillations. Thus, stimulus-dependent synchronous oscillations do not necessarily result in less information transfer, even when only the total spike count over a relatively long interval is considered.

Figure 7 about here

Hybrid Rate/Correlation Codes Work Best over Limited Numbers of Cells

Up to this point, we have only considered the effects of correlations between relatively small numbers of inputs. While this is consistent with the convergence ratios of sensory input onto neurons in the LGN and V1 [Reid, 1995 #1373; Usrey, 1999 #1836], hybrid rate/correlation codes might cease to be advantageous as the number of inputs is increased. To explore the behavior of a hybrid rate/correlation code as a function of the

number of input spike trains, we used a much larger stimulus that activated a 12×12 array of neurons. Oscillations between retinal ganglion cells increase markedly in response to the larger stimuli [Ishikane, 1999 #1934; Neuenschwander, 1999 #1766; Ariel, 1983 #538] and this was also true in our retinal model. However, because the maximum correlation strength between cat alpha cells has not been precisely determined for such stimuli, our results should be interpreted only as a rough estimate of the information that might be conveyed by a hybrid rate/correlation code as the number of inputs increases. As a function of the number of input spike trains, Δ reached a maximum for a relatively small number of correlated inputs, between 10 and 50 (fig. 8a). For integration times less than 5 msec, Δ remained greater than or equal to zero regardless of the number of inputs. For a threshold process employing longer integration times, correlations produced progressively poorer stimulus discrimination as the number of inputs increased.

As the number of inputs was raised, we increased the threshold so as to maintain the background detection rate as close to 1 Hz as possible, but without falling below 0.1 Hz. As the number of inputs increased, it was necessary to use a lower threshold for the Poisson control than for the correlated input, due to the small oscillations present in the background activity, in order to maintain the level of background suprathreshold events close to the target level of 1 Hz. We obtained similar results when the target level of background suprathreshold events was allowed to increase, but at smaller values of Δ on average. In order to maintain task difficulty as the number of inputs was increased, the duration of each discrimination trial was lowered from 400 to 100 msec.

The dependence of the hybrid code on the number of cells feeding into the detector is paralleled by the behavior of the Fano factor of the combined input (fig. 8b). When the number of inputs was small, the Fano factor was always less than one, regardless of the stimulus intensity. As the number of inputs increased, however, the Fano factor became much greater than one at all but the smallest stimulus intensities. This result is consistent with previous studies showing that synchrony becomes progressively more destructive of information encoded by the average firing rate as the number of neurons contributing to the

average increases [Mazurek, 2002 #1987; Shadlen, 1994 #1807; Shadlen, 1998 #1802]. When the number of inputs is very large, therefore, a hybrid rate/correlation code is only likely to be effective when the integration time is small enough to resolve the stimulus information encoded directly by the degree of synchrony itself.

Figure 8 about here

Discussion

Several influential studies of how firing synchrony affects the representation of information in neural ensembles focused solely on how spatial correlations impede the extraction of information from the averaged firing rates [Mazurek, 2002 #1987; Shadlen, 1994 #1807; Shadlen, 1998 #1802]. In particular, these studies emphasized how averaging over ensembles of similarly activated neurons only reduces the trial-to-trial variability in the estimated firing rate to the extent that the individual cells are uncorrelated. From this point of view, synchrony is an inevitable feature of densely interconnected networks and as such limit the effective size of neural ensembles to several tens of strongly correlated cells. However, these studies generally ignored the possibility that synchronous oscillations could themselves directly encode stimulus information, and thus compensate for the attendant loss of the information encoded by the averaged firing rate. Here, we have demonstrated that realistic synchronous oscillations can lead to improved information transmission under fairly general assumptions. In particular, synchrony, when evoked in an intensity-dependent manner, can significantly improve information transfer through threshold neurons functioning as coincidence detectors. In a complementary fashion, oscillations, when also proportional to intensity, mediate greater stimulus discrimination by making the total number of input spikes more reliable regardless of the integration time window employed. Thus, we found that the negative effects of firing synchrony described in previous studies could be overcome in two ways: 1) by encoding similar information in the degree of synchrony itself and 2) by making the spike trains more regular, via more pronounced oscillations, as they become more synchronous.

General principles of population encoding and decoding, especially those involving correlations between multiple neurons, are still very difficult to investigate experimentally. To reproduce the results of the present model, it would be necessary to record simultaneously from multiple input neurons and at least one common postsynaptic neuron and to manipulate the firing correlations between the input neurons either through direct multi-electrode stimulation or else via pharmacological techniques. While it is not clear that such biological

experiments are currently feasible, there does exist sufficient information to construct artificial spike trains possessing physiologically realistic correlations. Moreover, it is not necessary to possess a complete understanding of the physiological mechanisms that give rise to strong correlations in order to analyze how their information content might be extracted via a biologically plausible mechanism. Here, we have focused on signal detection using a threshold process that captures certain essential aspects of single neuron dynamics. By using a computational model, it was possible to compare the stimulus discrimination accomplished by a hybrid rate/correlation code with that mediated by a pure rate-code that produced, on average, the same number of spikes. Finally, by using a computational model, we were able to examine the effects of synchronous oscillations over a wide class of threshold detection processes, and in this way obtained insights into the physiological conditions necessary to utilize a hybrid rate/correlation code.

The main drawback of using a computational model is the need to ascertain whether the results are physiologically relevant. Fortunately, synchronous oscillations between retinal ganglion cells have been sufficiently well characterized so as to impose tight constraints on artificially generated spike trains. Where possible to verify, the model retina appeared to favor a rate code over a correlation code. This was the case for the narrow bar stimulus employed in our first several experiments. The maximum synchrony between the model spike trains, measured relative to chance and averaged over all cells responding to the bar, was somewhat less than the levels of synchrony between widely separated cells recorded in the cat retina in response to similar stimuli [Neuenschwander, 1996 #1331]. Likewise, the maximum sustained increase in firing rate in our artificial spike trains was probably exaggerated due to the lack of several known adaptation mechanisms in the model, particularly those present in the outer retina. Our results were also useful for illustrating general phenomena, such as how hybrid rate/correlation codes might scale as the number of cells conveying redundant information increased. The effectiveness of a hybrid rate/correlation code was found to involve a trade off between the loss of information in the averaged firing rates due to synchrony and a corresponding gain in information due to the

information encoded by the synchronous oscillations themselves. At longer integration times that were insensitive to synchronous input, stimulus-dependent oscillations still contributed to improved performance on intensity discrimination tasks by causing spike counts to become more regular. As the number of neurons increased, the loss of rate-coded information due to synchrony became more severe.

Our results are likely to be helpful for interpreting synchronous oscillations in the retina and elsewhere in the CNS. Gamma band oscillations are ubiquitous in the vertebrate retina, having been measured extracellularly in cats [Laufer, 1967 #1964; Steinberg, 1966 #1963; Neuenschwander, 1996 #1331; Neuenschwander, 1999 #1766], rabbits [Ariel, 1983 #538], frogs [Ishikane, 1999 #1934] and mudpuppy [Wachtmeister, 1978 #1962], as well as in the ERGs of humans [De Carli, 2001 #1930; Wachtmeister, 1998 #1793] and primates [Frishman, 2000 #1931]. The conservation of retinal oscillations across such a broad range of vertebrate species suggests they may be important for visual function. Synchronous oscillations have also been recorded elsewhere in the mammalian nervous system, including visual [Gray, 1989 #1703; Gray, 1989 #1704; Kreiter, 1996 #1730; Livingstone, 1996 #1538] and sensorimotor [Murthy, 1996 #1935; Murthy, 1996 #1936] cortex and the hippocampus [Traub, 1996 #1648]. Numerous interpretations of the information processing function accomplished by synchronous oscillations have been suggested [Engel, 2001 #1943; Fries, 2001 #1986; Fries, 2002 #1940; Singer, 1995 #1367; Srinivasan, 1999 #1757]. Here, we note that since synchronous oscillations are widely present throughout the brain; the nervous system might as well make use them. If correlations are indeed an unavoidable consequence of neural connectivity, our results suggests that the loss of information encoded by the average firing rate could be mitigated by employing a hybrid rate/correlation code. On the other hand, correlations may also convey information that is not well represented in average firing rate. Our results suggest that the brain might exploit synchronous oscillations to represent additional types of information without sacrificing total information content.

Several of our results were predicated on an analysis of the correlations present during the plateau portion of the response, using a discrimination window of 200 msec. While the

majority of stimulus information is probably transmitted during the first 50 msec encompassing the response peak, additional information is likely transmitted by later components as well. The mean inter-saccade interval for primates during a free viewing visual search task is approximately 200 msec [Mazer, 2003 #2261], while optimal temporal frequencies for neurons in areas 17/18 in the cat visual cortex are typically in the range of a few Hz [Movshon, 1978 #2263], implying sustained activations on the order of several hundred msec. Moreover, stimulus-related increases in gamma band power have recently been shown to persist for at least 1 sec during free viewing conditions [Salazar, 2004 #2262]. Thus, the 200 msec time window used in several of our experiments is broadly consistent with relevant physiological time scales. In addition, we found that stimulus-dependent synchronous oscillations mediated improved discrimination even for analysis windows confined to the first 50 msec following stimulus onset, despite the presence of strong additional correlations due to the response peak itself in both the model and control data.

It has been argued that information transfer is maximized when retinal outputs are uncorrelated [Srinivasan, 1982 #1829] and recent studies indicate that the relatively strong firing correlations sometimes observed between neighboring ganglion cells convey little additional information about natural scenes [Nirenberg, 2001 #1881]. Our results are not in conflict with these findings. In the present analysis, retinal inputs were pooled into a single variable, roughly corresponding to a postsynaptic membrane potential. The only benefit of statistical independence was therefore a possible improvement in the signal to noise of the pooled input. Our results, however, suggest that any reduction in the signal to noise of the pooled input due to synchrony could be outweighed by the stimulus information encoded by the synchrony itself, as well as by the attendant reduction in overall spike count variability resulting from the presence of stimulus-dependent oscillations. Moreover, information theory, by itself, does not address how stimulus properties are extracted by target neurons. While formal mathematical measures predict that ganglion cells convey more information when their activity is uncorrelated, experimental and theoretical evidence suggests that synchronous inputs are particularly salient [Usrey, 2000 #1833; Alonso, 1996 #1370;

Kenyon, 1990 #1375], and our current results demonstrate that information encoded by firing correlations between model ganglion cells can be efficiently extracted by threshold neurons under fairly general assumptions. Simultaneous recordings in cat from the retina, the LGN, and from area 18 of the visual cortex indicate that synchrony between retinal ganglion cells can be propagated to higher levels in the visual system [Castelo-Branco, 1998 #1497]. It may therefore be necessary to consider the impact of correlated activity in order to fully account for its role in information processing. Finally, we note that using firing correlations to encode local stimulus properties does not preclude additional encoding functions that have been suggested [Meister, 1999 #1666].

Correspondence and requests for materials should be addressed to G.K. (e-mail: gkenyon@lanl.gov).

Figure Captions

Fig. 1. The retinal model contained five cells types: bipolar (BP) cells, small (SA), large (LA) and poly-axonal (PA) amacrine cells, and alpha ganglion (GC) cells, arranged as a 32x32 square mosaic with wrap-around boundary conditions. Conceptually, connections could be organized into 3 categories. a) Feedforward and feedback inhibition. Excitatory synapses from BPs were balanced by a combination of reciprocal synapses and direct inhibition of the GCs, primarily mediated by the non-spiking amacrine cell types. b) Serial inhibition. The three amacrine cell types regulated each other through a negative feedback loop. c) Resonance circuit. The PAs were excited locally via electrical synapses with GCs and whose axons mediated wide field inhibition to all cell types, but most strongly onto the GCs. Not all connections are shown. Explanation of symbols: Excitation (triangles), inhibition (circles), gap junctions (resistors).

Fig 2. Artificial spike trains generated by the retinal model. a) A column of twelve model ganglion cells was stimulated by a narrow bar (white rectangle, intensity = ½, stimulus dimensions 2×12 GC receptive field diameters). Circles indicate GC receptive field diameter. b) Multiunit Peri-Stimulus-Time-Histogram (PSTH) obtained by averaging the individual PSTHs over all ganglion cells activated by the stimulus (bin width, 1 msec). The solid line at the bottom of the panel indicates the stimulus duration (600 msec). Vertical ticks denote the peak and plateau portions of the response, respectively. c) Multiunit Cross-Correlation Histogram (CCH), obtained by combining individual CCHs from all distinct pairs of ganglion cells activated by the stimulus measured during the plateau portion of the response (solid black lines). Correlations expressed as a fractional change from the expected synchrony due to chance (dimensionless units). Shift predictors (dashed gray lines) obtained by recomputing the multiunit CCH using spike trains from different stimulus trials. d) Multiunit CCH recorded measured during the response peak (same organization as panel c). Correlations were larger during the response peak, as was the shift predictor, due to the strong stimulus-locking of the high frequency oscillations.

Fig. 3. Firing correlations are modulated over a greater dynamic range than firing rate as a function of stimulus intensity. The stimulus was again a narrow bar covering twelve ganglion cells. a) Intensity series formed by the multiunit PSTHs of ganglion cells activated by the stimulus. The intensity, in $\log_2$ units, is indicated at the upper right of each histogram (bin width = 10 msec). b) Intensity series formed by the multiunit CCHs between all pairs of ganglion cells activated by the stimulus, relative to the baseline level of synchrony. Firing correlations during the plateau response are strongly modulated by stimulus intensity. c) Intensity series formed by the multiunit CCHs recorded during the response peak (same organization as panel b). d,e) Fractional change from baseline in synchrony (black lines, circles) as a function of stimulus intensity, compared to the fractional change in firing rate (gray lines, squares), recorded during the d) plateau and e) peak portions of the response. In both cases, synchrony was modulated over a greater dynamic range than the firing rate.

Fig. 4. Firing correlations during the plateau portion of the response allow improved discrimination of stimulus intensity compared to independent Poisson input. a) Example of the threshold detection process. Stimuli consisted of a narrow bar presented at various intensities (same stimulus as in fig. 3). Ganglion cell input to the threshold detector during the plateau portion of the response is shown on the left and equivalent Poisson input on the right. The baseline activity of the detector in the absence of stimulation (top row) is very low. A stimulus with intensity = -1 in $\log_2$ units (bottom row) produced strong correlations between ganglion cells, resulting in a relatively large number of suprathreshold events. Dashed line: detector threshold. Dotted lines: average input ± std. dev. Summation window, 2 msec. b) Normalized probability distributions giving, for each stimulus intensity, the relative number of suprathreshold events during the 200 msec analysis interval. Left: Retinal model. Right: Poisson control. The distribution of suprathreshold events produced by the retinal model were more separable. c) Fraction of correct intensity classifications by an ideal observer. The ideal observer chose between two equally likely stimulus intensities based on the number of detector events on each trial. Input to the detector came from either the retinal model (solid black lines) or from Poisson controls (dashed gray lines). Each point represents

the fraction of trials on which the intensity indicated by the abscissa was correctly distinguished from a lower intensity, denoted by the intersection of each line with the x-axis. Error bars computed assuming binary statistics for the overlap between each pair of distributions (omitted from Poisson controls for clarity).

Figure 5. Firing correlations during the response peak allow improved discrimination of stimulus intensity. Same organization as in fig. 4. To remove the high-frequency stimulus-locked oscillations present in the response peak (panel a, bottom left), the multiunit PSTH was smoothed by replacing the measured firing rate in each bin by the average of the surrounding nine bins (panel a, bottom right). Firing correlations during the response peak allowed for improved discrimination of stimulus intensity compared to the case where correlations were due solely to the response transient itself (panels b and c).

Figure 6. Different mechanisms for generating synchronous oscillations yield similar stimulus discrimination. a) Left Column: Multiunit CCHs generated by the retinal model in response to stimuli of varying intensity, starting with no stimulus (bottom row) and ranging from –6 to –1 ($\log_2$ units, top row). Multiunit CCHs expressed as a fraction of asymptotic baseline level. Right Column: Multiunit CCHs produce by an equal number of Poisson event generators whose rates were modulated by a source of common oscillatory input (same organization as in left column). Both sets of artificial spike trains exhibited similar correlations as measured by the multiunit CCH but may have differed at the single trial level. b) Despite employing very different mechanisms, stimulus discrimination was similar for the two models (same parameters as used in figure 4).

Figure 7. Hybrid rate/correlation codes yield superior or equivalent stimulus discrimination for a broad class of threshold detectors. a) Δ, the difference in the percentage of successfully classified trials using a hybrid code as opposed to independent Poisson controls, plotted as a function of temporal integration window. Δ declined with increasing integration time, but did not become strongly negative. Individual points are for different thresholds. b) Δ vs.

detection threshold for different integration times. Same data as in previous panel. The increase in Δ with threshold declines progressively at longer integration times. c) Fano factors (variance/mean of the total input spike count) plotted vs. stimulus intensity. Solid line+circles: Hybrid code. Dashed line+squares: Poisson control. Dotted line+triangles: Modified Poisson process in which the separate generators were synchronized by an amount proportional to the stimulus intensity in the absence of oscillations. To account for the effects of a refractory period on the Fano factor, no more than one spike was allowed to occur in any 10 msec time bin. The Fano factor for the hybrid code remained less than one due to the presence of high frequency oscillations, while synchrony alone produced a large increase in variability relative to independent Poisson inputs.

Figure 8. Hybrid rate/correlation codes work optimally for limited numbers of input neurons. These experiments used a 12×12 uniform spot that produced larger synchronous oscillations than did a narrow bar. a) Δ vs. the number of input cells plotted for several different integration times. Δ reached a peak at between 10-50 inputs and remained positive as the number of inputs increased as long as the integration time was small, but Δ becomes negative for the same number of inputs if the integration time was too long to resolve synchronous inputs. b) Fano factor vs. number of inputs plotted for several different intensities ($\log_2$ units). At high intensities, which produce strong synchronous oscillations, the Fano factor increased sharply as a function of the number of inputs, thus accounting for the poorer performance of the hybrid code in this regime.

Table 1: Cellular parameters.

	τ	b	n×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 (msec); b: bias; n×n: array size; d: cutoff radius, σ: Gaussian radius (see eq. 5). [a]Inner radius/outer radius.

Table 2: Synaptic weights.

	BP	SA	LA	PA	GC
BP	*	-0.375[b]-	3.0b	-3.0b/-15.0c	*
SA	3.0[b]	*	-3.0[b]	0.0b/-15.0[c]	*
LA	3.0[b]	*	0.25[a]	-3.0a/-15.0[c]	*
PA	0.75[b]	-0.75b	0.25[a]	0.25a/-45.0[c]	0.25[a,d]
GC	9.0[b]	-4.5b	-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. 5). 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%.

References:

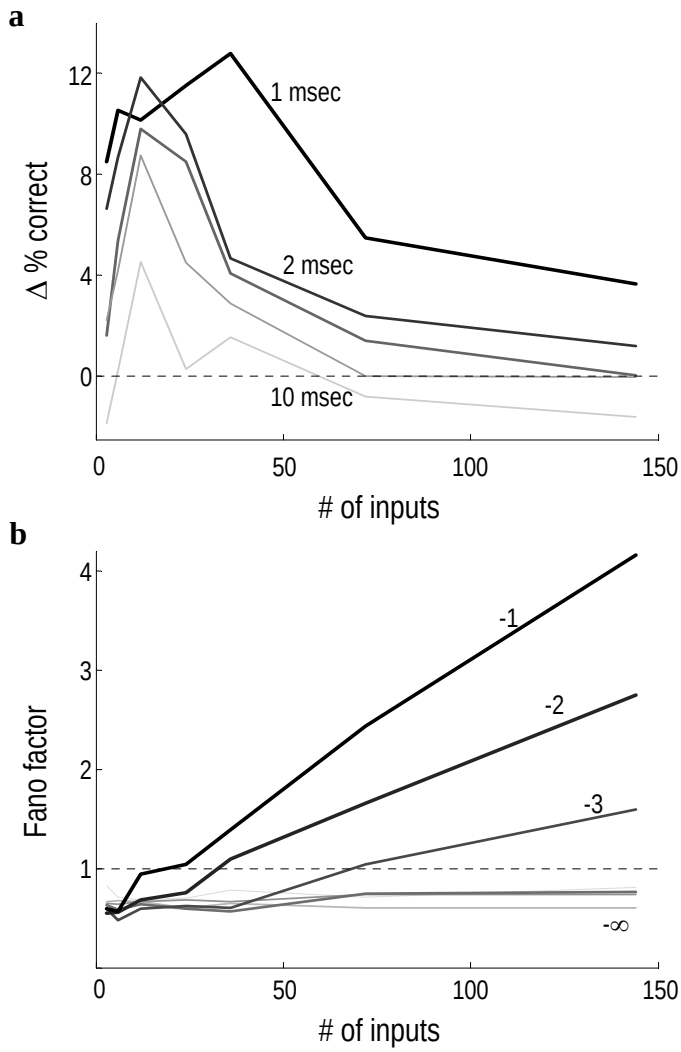
a
12
8
4
0
Δ % correct
1 msec
2 msec
10 msec
0
50
100
150
# of inputs
b
4
3
2
1
Fano factor
-1
-2
-3
-∞
0
50
100
150
# of inputs

Figure 8