

Correlated Firing in Macaque Visual Area MT: Time Scales and Relationship to Behavior

Wyeth Bair,^{1,2} Ehud Zohary,³ and William T. Newsome^{1,4}

¹Howard Hughes Medical Institute (HHMI), ²Center for Neural Science, New York University, New York, New York 10003, ³Department of Neurobiology, Institute of Life Science, Hebrew University, Jerusalem, 91904, Israel, and ⁴Department of Neurobiology, Stanford University School of Medicine, Stanford, California 94305

We studied the simultaneous activity of pairs of neurons recorded with a single electrode in visual cortical area MT while monkeys performed a direction discrimination task. Previously, we reported the strength of interneuronal correlation of spike count on the time scale of the behavioral epoch (2 sec) and noted its potential impact on signal pooling (Zohary et al., 1994). We have now examined correlation at longer and shorter time scales and found that pair-wise cross-correlation was predominantly short term (10–100 msec). Narrow, central peaks in the spike train cross-correlograms were largely responsible for correlated spike counts on the time scale of the behavioral epoch. Longer-term (many seconds to minutes) changes in the responsiveness of single neurons were observed in auto-correlations; however, these slow changes in time were on average uncorrelated between neurons. Knowledge of the limited time

scale of correlation allowed the derivation of a more efficient metric for spike count correlation based on spike timing information, and it also revealed a potential relative advantage of larger neuronal pools for shorter integration times. Finally, correlation did not depend on the presence of the visual stimulus or the behavioral choice of the animal. It varied little with stimulus condition but was stronger between neurons with similar direction tuning curves. Taken together, our results strengthen the view that common input, common stimulus selectivity, and common noise are tightly linked in functioning cortical circuits.

Key words: Area MT/V5; cross-correlation; neuronal pooling; visual motion; extrastriate cortex; synchrony; stimulus-locked modulation; noise correlation; visual cortex

A fundamental problem in sensory neuroscience is to understand how psychophysical performance is related to the signaling capacities of single sensory neurons. It is now widely recognized that no satisfactory solution to this problem can be achieved in the absence of detailed knowledge concerning correlated firing within the pool of sensory neurons contributing to a particular psychophysical judgment (Johnson et al., 1973; Johnson, 1980; van Kan et al., 1985; Britten et al., 1992; Gawne and Richmond, 1993; Zohary et al., 1994; Geisler and Albrecht, 1997; Parker and Newsome, 1998). For example, combining signals across a pool of neurons can generate superior psychophysical sensitivity if the noise carried by individual members of the pool is averaged out. This benefit of pooling is only achievable, however, to the extent that the noise carried by individual neurons is independent (uncorrelated); noise that is common to the entire pool cannot be averaged out. In general, the effect of correlated noise depends on how signals are combined, and although correlation may either aid or hinder noise removal (Johnson, 1980; Abbott and Dayan, 1999; Panzeri et al., 1999), its impact on the amount of information conveyed by a pool of neurons may be profound. Thus, empirical analysis of correlated firing is central to a quantitative

understanding of the relationship between physiological responses and psychophysical judgments.

Extrastriate visual area MT is ideal for investigating pools of sensory neurons that underlie psychophysical performance. MT contains a preponderance of directionally selective neurons (Zeki, 1974; Maunsell and Van Essen, 1983; Albright et al., 1984), the activity of which has been linked compellingly to the psychophysical discrimination of direction in stochastic motion stimuli (Newsome et al., 1989; Britten et al., 1992; Salzman et al., 1992; Murasugi et al., 1993; Salzman and Newsome, 1994). In a previous study, therefore, we measured correlated firing in MT and found that spike counts from adjacent neurons were noisy and only weakly correlated but that even this small amount of correlated noise placed substantial limits on the benefits of signal averaging across a pool (Zohary et al., 1994). Subsequently, Shadlen and colleagues (1996) incorporated these insights into a computational model of the relationship between the activity of MT neurons and psychophysical judgments of motion direction.

In the present study, our primary goals were to examine the time scale at which correlation arises: in particular, to relate spike count correlation to spike timing correlation and examine the dependence of correlated firing on stimulus and behavioral parameters. Our most intriguing finding is that trial-to-trial correlations in spike count, measured over trials of 2 sec duration, are produced largely by the same mechanisms that generate peaks in the spike train cross-correlogram (CCG) on a time scale of a few tens of milliseconds. For a given pair of MT neurons, a quantitative measurement based on the CCG peak predicts with fair accuracy the level of correlation calculated from spike counts over the full trial length. Furthermore, the CCG-based measure is

Received Sept. 22, 2000; revised Dec. 5, 2000; accepted Dec. 12, 2000.

W.B. is supported by Howard Hughes Medical Institute (HHMI). Part of this work was funded by the L. A. Hansen Fellowship to W.B. while in the lab of Christof Koch at Caltech. W.T.N. is an investigator of HHMI. We thank Michael N. Shadlen, Carlos Brody, J. Anthony Movshon, and Christof Koch for suggestions and helpful discussion that has guided the course of this work, and we owe additional thanks to M. N. Shadlen and C. Brody for detailed comments on this manuscript.

Correspondence should be addressed to Wyeth Bair, Howard Hughes Medical Institute, Center for Neural Science, New York University, 4 Washington Place, Room 809, New York, NY 10003. E-mail: wyeth@cns.nyu.edu.

Copyright © 2001 Society for Neuroscience 0270-6474/01/211676-22\$15.00/0

substantially more reliable than the measure based on spike count correlation. The spike train CCG is typically used as a qualitative indicator of functional connectivity among neurons. In contrast, our results suggest that the spike train CCG can provide quantitative measures of neuronal correlation that are of considerable interest for models that seek to reconcile neuronal and psychophysical performance.

Some of these results have been published previously in abstract form (Bair et al., 1996, 1999).

MATERIALS AND METHODS

Subjects, surgery, and daily routine. The experiments were performed on three adult rhesus monkeys weighing between 7 and 9 kg (*Macaca mulatta*, two males and one female). Before the experiments, each monkey was surgically implanted with a device for stabilizing head position (Evarts, 1968), a scleral search coil for measuring eye position (Judge et al., 1980), and a recording cylinder that allowed microelectrode access to cortex within the occipital lobe. All surgical procedures were performed under aseptic conditions with halothane anesthesia. After recovery from surgery, each animal engaged in daily training or experimental sessions of 2–6 hr duration. Behavioral control was accomplished by operant conditioning techniques using fluids as a positive reward; fluid intake was therefore restricted during periods of training or electrophysiological recording. The diet was supplemented with moist monkey treats, fruits, and nuts. The animals were maintained in accordance with guidelines set by the U.S. Department of Health and Human Services (NIH) *Guide for the Care and Use of Laboratory Animals*.

Visual stimuli. The visual stimuli used in this study were a set of dynamic random dot patterns in which a unidirectional motion signal was interspersed among random motion noise. The stimulus set has been described extensively in previous publications (Britten et al., 1992), and we simply summarize its essential features here.

Dynamic random dots were plotted sequentially on the face of a CRT screen at a high rate (6.67 kHz). After 45 msec, a dot was either displaced in a specified direction (coherent motion) or replaced by another dot at a random location on the screen (noise). In one extreme form of the display, all dots were positioned randomly so that the display was pure noise. In this form, which we term 0% coherence, the display contained many local motion events (caused by fortuitous pairings of the dots in space and time) but on average no net motion in any direction. At the other extreme (100% coherence), all dots were displaced uniformly so that the display contained noise-free motion in a specified direction. Our software permitted us to create any stimulus intermediate between these two extremes by specifying the percentage of dots that carried the “coherent” motion signal. The percentage of dots engaged in coherent motion governed the strength of the motion signal without affecting the overall luminance, contrast, or average spatial and temporal structure of the stimulus. When a psychophysical subject was asked to discriminate the direction of motion in such displays, the difficulty of the discrimination was related directly to the percentage of dots in coherent motion.

In early experiments (monkey E), the stimuli were generated by a PDP 11/73 computer and displayed on a large, electrostatic deflection oscilloscope via a high speed DMA digital-to-analog converter. In later experiments (monkeys R and K), the stimuli were created by means of an IBM 386 equipped with a dedicated graphics board (SGT Pepper no. 9). These stimuli were displayed on a raster scan CRT monitor with a 60 Hz refresh rate. In all experiments, the display monitor was positioned 57 cm in front of the monkey.

A critical distinction must be made between two different methods of presenting repeated stimuli for a particular condition (e.g., a 6.4% coherence, upward stimulus). Our standard method used a new random number sequence for each repeat, resulting in what we will refer to as “ensemble stimuli,” which differ in detail but have on average the prescribed motion coherence. As a control for the effect of random stimulus variation on neuronal responses, we recorded from four pairs using repeated presentations of stimuli generated with exactly the same sequence of random numbers. We will refer to the identical stimulus repeats used by this method as “replicate stimuli.”

Behavioral paradigms and selection of visual stimuli. We used two behavioral paradigms in this study: a fixation task and a discrimination task. In the fixation task, the monkey was required only to maintain its eye position within an electronically defined window around the fixation point for 2–4 sec. The monkey received a liquid reward on successful

completion of each trial. In most experiments, the window permitted eye movements up to 1.5° away from the fixation point, but in practice, the monkeys usually held their eye position within 0.5° of the fixation point.

The monkeys performed the fixation task during the initial search for well isolated pairs of neurons, during mapping of receptive fields, and during quantitative measurement of the direction tuning properties of the neurons. Receptive field boundaries were mapped qualitatively for each neuron of the pair, and the stimulus aperture was positioned to include both receptive fields. The receptive fields typically overlapped substantially, so that the stimulus aperture only engaged a small portion of the surround of either receptive field. The optimal speed was estimated qualitatively for each neuron, and subsequent experiments were conducted using a motion speed intermediate between the two optima. To measure a direction tuning curve, a 100% coherence dot pattern was presented in eight different directions of motion equally spaced around the clock at 45° intervals. The different directions were presented in a pseudorandom sequence until 10–20 repetitions were completed for each direction.

The two direction tuning curves were used to assign a “preferred-null” axis of motion for use during the discrimination task (below). Ostensibly, the preferred-null axis was the axis of maximal directionality for the two neurons; motion in opposite directions along the axis should yield a maximal difference in responsiveness. In practice the axis chosen was usually a compromise between the preferred directions of the two neurons measured individually. Most pairs of neurons had similar preferred directions, and the compromise therefore resulted in a near-optimal axis for both. On occasion we recorded from pairs of neurons with preferred directions that were nearly opposite each other. In this case again, the choice of directional axis was easy because the signs of the two response were simply reversed along the same axis. Occasionally, however, the preferred directions of the two neurons were nearly orthogonal to each other, or one of the neurons was not directional at all. In such cases, we chose the preferred-null axis and the speed of the motion signal to match the preferences of the more responsive, directional neuron. On the whole, therefore, most neurons were studied during the discrimination task with stimuli that matched their physiological properties reasonably well. For a few neurons, the stimuli were substantially suboptimal.

In the discrimination task, the monkey performed a two-alternative, forced-choice discrimination of motion direction. This task has been used extensively in our laboratory and is described in detail in previous publications (Britten et al., 1992). On each trial a random dot stimulus was presented for 2 sec within the aperture covering both receptive fields. The direction of the coherent motion signal was varied randomly from trial to trial between the preferred direction of the neurons under study and the direction 180° opposite (the “null” direction); the monkey’s task was to discriminate correctly the direction of motion. The strength of the motion signal was varied among a range of coherence levels that spanned psychophysical threshold. A minimum of 15 repetitions was obtained for each stimulus condition (i.e., each combination of direction and coherence), and all conditions were presented in pseudorandom order. We will refer to the neuronal data from these experiments as “coherence series data” to distinguish them from the direction tuning data.

Each trial began with the appearance of the fixation point. After the monkey achieved fixation and held its gaze within the fixation window for 300 msec, the visual stimulus was presented as described above. The monkey was required to hold its gaze on the fixation point during stimulus presentation so that the stimulus remained well positioned on the receptive fields of the two neurons. At the end of the 2 sec display interval, the random dot pattern and the fixation point disappeared, and two small visual targets appeared, one corresponding to each of the two possible directions of coherent motion. The monkey made a saccadic eye movement to one of the two targets to indicate the direction of motion perceived in the visual stimulus. Eye movements were measured continuously with the scleral search coil technique, permitting the computer to register correct and incorrect choices. Correct choices were followed by a liquid reward; incorrect choices were followed by a brief time-out period. On 0% coherence trials, the monkey was rewarded randomly with a probability of 0.5 because there was no “correct” answer on these trials. If the monkey broke fixation prematurely during a trial, the trial was aborted, the data were discarded, and a time-out period ensued.

Electrophysiological recording and spike sorting. Electrophysiological recordings were made with tungsten microelectrodes inserted into the cortex through a transdural guide tube (electrode impedance = 0.5–2.0 MΩ at 1 kHz) (Micro Probe, Potomac, MD). The guide tube was held

rigidly in a stable coordinate system by a plastic grid inside the recording cylinder (Crist et al., 1988). We recorded through any particular guide tube for several consecutive days.

The signal from the microelectrode was amplified and bandpass-filtered (0.5–10 kHz), and action potentials from multiple single neurons were discriminated using an on-line spike sorting system that was developed originally in the laboratory of Dr. Moshe Abeles (Hebrew University, Jerusalem) and was commercially available from Alpha Omega Engineering (Nazareth, Israel). The filtered microelectrode signal was continuously sampled at a rate of 14 kHz by a digital signal processing system housed in an IBM 386 platform. (The apparent discrepancy between the 7 kHz cutoff frequency, implied by 14 kHz sampling, and the 10 kHz cutoff of our bandpass filter was not a limiting factor, because in practice the amplitude of the noise from 7 to 10 kHz was small relative to the amplitude of all well isolated action potentials.) The computer software provided a user interface to the spike sorting hardware and included graphics displays of voltage waveforms, spike templates, and distributions of matching errors (below). Spikes were discriminated on-line using an eight-point template-matching algorithm (Wörgötter et al., 1986). Each time the voltage exceeded a threshold level, an eight-point voltage sample was acquired and compared with the predefined templates that characterized the waveform of each recorded neuron. If the root-mean-square error (RMSE) of the match between the signal waveform and one of the templates was below a criterion value, an action potential was registered for that neuron. A template was defined by the software to minimize the RMSE of the match to the template across a sample of 100 action potentials accepted by the experimenter as belonging to a specific neuron.

The quality of unit isolation was determined by the separation of the templates from each other and from the noise. Excellent separation of the templates from each other was necessary to prevent “cross-talk” between the two waveforms. Occasional misclassification of the two action potentials could result in artifactual correlations that would be deleterious to certain analyses. The only substantive insurance against cross-talk were the rigor and attentiveness of the experimenter—both in selection of pairs for study and in maintaining quality of isolation during the experiment. We attempted to be exceedingly rigorous in selecting pairs for study, rejecting all candidates except those with waveforms that were strikingly distinct from each other. Similarly, we attempted to be unusually conservative in on-line assessment of the quality of isolation. If either waveform began to deteriorate, creating any doubt about isolation, we ceased recording until the waveforms could be restored.

Separation of the two templates from the noise could be achieved more objectively. For each template, the software compiled and displayed a frequency histogram of RMSE values resulting from comparison of each triggered waveform with that template. Excellent separation of the template from the noise corresponded to a bimodal histogram of RMSE values. A peak at low RMSE values corresponded to action potentials from the neuron that defined the template; a larger peak at high RMSE values reflected the substantial mismatch between noise waveforms and the template. We insisted that both modes be visible and well separated from each other in the RMSE histogram. The criterion RMSE value for accepting an action potential as corresponding to a particular template was set at the local minimum in the bimodal RMSE distribution. This ensured a reasonable balance between minimizing noise contamination and minimizing false negative matches to the template. We rejected recordings for which the error distributions were judged by eye to overlap in a manner that would produce more than ~5% false positives, but we estimate that the contamination rate was typically lower because the peaks in the bimodal error distribution often showed no sign of overlap after collecting hundreds of spikes. Admitting a small percentage of spikes from other neurons to one or the other template should have negligible effect on estimates of pair-wise interneuronal correlation because of the modest to weak correlations typical between cortical neurons.

Obviously, our technique of multi-unit recording with a single electrode cannot detect simultaneous spikes because the two waveforms superimpose, resulting in a poor match to either template. Because the primary lobes of the action potential waveforms were generally ≤ 0.5 msec in duration (Mountcastle et al., 1969; Funahashi and Inoue, 2000), this limitation only resulted in an underestimation of spikes that were synchronous to within 1 msec [for example, see Gawne and Richmond (1993); Funahashi and Inoue (2000)]. For two neurons with uncorrelated activity firing at rates < 100 spikes per second, the probability of spike synchrony at the millisecond time scale is $< 0.1^2$, which is reasonably

uncommon. However, the pairs of neurons studied here often have peaks in their CCGs at time zero (see Fig. 5B), and therefore the probability of simultaneous firing may be many times greater. This problem can be compounded for cells that fire bursts during which firing rates may reach 300–500 spikes per second. However, multi-electrode cross-correlation studies in monkey and cat suggest that CCG peaks in visual cortex are typically broader than 1–2 msec (Ts'o et al., 1986; Krüger and Aiple, 1988; Ts'o and Gilbert, 1988; Cardoso de Oliveira et al., 1997). The available evidence suggests that the vast majority of CCGs do not have sudden discontinuities on the time scale of 1 msec at the origin and that peaks of width 1 msec, when they exist, are weak and could not account for a substantial fraction of the strength of interneuronal correlation commonly observed in visual cortex. Therefore, we approximate the CCG value at time zero using values at neighboring time lags, as described later.

Analysis of direction selectivity. To assess neuronal direction selectivity, we determined which of two different models could better match the direction tuning curves. The first model assumed that the neuron was not direction selective and that response variation across direction was caused simply by sampling noise. It therefore predicted that the level of activity was essentially invariant with direction and was best estimated as the mean of the responses to all directions. The second model assumed that the neuron was in fact direction selective with a Gaussian distribution of responses centered on the optimal direction of motion. This distribution had four free parameters: the optimal direction of motion, the maximal response rate, the bandwidth of the Gaussian function, and the baseline response (the spontaneous firing rate). We performed maximum likelihood fits to the two separate, nested models under the assumption of normal errors. The likelihoods (L) obtained from these computations were transformed by:

$$l = -2 \ln \frac{L(\text{data}|\text{Gaussian fit})}{L(\text{data}|\text{uniform fit})}, \quad (1)$$

such that l is distributed as χ^2 with three degrees of freedom (Hoel et al., 1971). If l was below the criterion value ($p = 0.05$), we concluded that the direction tuning function of the neuron was better described by a Gaussian fit than by a constant response independent of direction. We considered these neurons to be direction selective, and the quantitative analyses in this paper used optimal directions and bandwidths obtained from the Gaussian fit to the tuning curve of each neuron. We will use the notation Δ_{PD} to refer to the difference (in degrees) between the preferred directions of neurons within a pair.

Analysis of psychophysical data. Psychophysical data from the discrimination experiments were compiled into psychometric functions depicting the proportion of correct decisions as a function of the strength of the motion signal (in % coherence). We used a maximum likelihood method (Watson, 1979) to fit these data with sigmoidal functions of the form:

$$p(c) = d - (d - 0.5)e^{-(c/a)^b}, \quad (2)$$

where p is the probability of a correct decision, c is coherence, a is the coherence level that supports threshold performance (82% correct), b is the slope of the sigmoidal function, and d is the asymptotic performance for strong motion signals (expressed as proportion of correct decisions). The threshold parameter, a , and the slope parameter, b , provide a succinct description of the psychophysical data.

Equation 2, derived from the integral of a Weibull distribution (Quick, 1974), provided acceptable fits to the bulk of our psychophysical data. Thirty-four of 46 psychometric functions in our data set were well fit [likelihood ratio test, $p > 0.05$; see the Appendix of Watson (1979)] when the asymptotic performance, d , was constrained to be unity, and the remaining functions were well fit by allowing d to vary. The non-unity asymptote in the latter 12 experiments reflected the monkey's occasional errors at the highest coherence levels.

Analysis of neural thresholds. We measured neural thresholds to the stochastic motion stimuli in a manner that permitted direct comparison with psychophysical thresholds [see Britten et al. (1992) for a detailed description]. For each neuron, we first compiled for each motion coherence a frequency histogram of responses to preferred direction motion and a separate histogram of responses to null direction motion. We considered a “response” to be the total number of spikes generated by the neuron during the 2 sec stimulus. For very strong (high coherence) motion signals, these “preferred” and “null” response distributions were typically non-overlapping because most of our neurons were highly directional. At these coherence levels, the direction of motion could be

determined unambiguously on any given trial simply by monitoring the response of the neuron. For very weak motion signals, however, the preferred and null response distributions overlapped almost completely, so judgments of motion direction based on the responses of the neuron would be at chance. Intermediate coherence levels resulted in partial overlap between the two response distributions, leading to intermediate levels of performance.

Following these intuitions, we used a method based on signal detection theory (Green and Swets, 1966; Britten et al., 1992) to compute the performance expected of an ideal observer who based judgments of motion direction on the measured neuronal responses. For each neuron, this calculation was performed for each coherence level (typically six non-zero levels, but as many as eight, and the results were compiled into a neurometric function that plotted expected performance (in % correct decisions) as a function of coherence. A sigmoidal curve was fitted to the data using Equation 2, and the threshold and slope parameters were extracted as described in the preceding section. These parameters describe the sensitivity of a single neuron to the motion signals in our displays in a manner that can be compared directly with the psychophysical sensitivity measured on the same trials. Equation 2 described our neurometric data well; the fits were acceptable (likelihood ratio test, $p > 0.05$) for all 83 of the neurons comprising the 46 pairs with valid psychophysical data.

Assessment of correlated activity. We analyzed two main types of correlation between the responses from each pair of neurons: signal correlation and noise correlation (Gawne and Richmond, 1993; Gawne et al., 1996; Lee et al., 1998). Signal correlation, designated r_{signal} , refers to the common modulation in a set of paired mean responses associated with multiple stimulus conditions. For our purposes, it is simply the correlation coefficient computed for the mean spike rates from a pair of direction tuning curves. Noise correlation, r_{noise} , refers to common trial-to-trial fluctuations around the mean response for a single stimulus condition, and its estimation and interpretation occupy the bulk of this paper. The dichotomy implied by the names “signal” and “noise” correlation is somewhat unfortunate because apparently noisy variations in spike rate may carry information about neural signals that we simply cannot access. However, we will adhere to these terms for the sake of precedent. The traditional measure of noise correlation is the interneuronal correlation coefficient (van Kan et al., 1985; Bach and Krüger, 1986; Gawne and Richmond, 1993; Zohary et al., 1994; Gawne et al., 1996; Lee et al., 1998), which measures correlation at a fixed time scale and temporal relationship, i.e., the simultaneous trial. We will describe two new methods for quantifying noise correlation, one at time scales greater than and equal to the single trial that generalizes the interneuronal correlation coefficient to non-simultaneous trials (below) and another at the scale of milliseconds that is derived from spike train correlograms (Appendix A). Table 1 provides a unified reference to all of our notation regarding correlation.

The trial cross-covariance. The interneuronal correlation coefficient is traditionally computed for the spike counts N_1 and N_2 of neurons 1 and 2, respectively, according to:

$$r_{\text{SC}} = \frac{E[N_1 N_2] - E N_1 E N_2}{\sigma_{N_1} \sigma_{N_2}}, \quad (3)$$

where E is expected value and σ is the SD computed across all repetitions of a particular stimulus. However, an experiment yields several sets of paired spike counts (one set for each stimulus condition), and rather than applying Equation 3 separately to each set, the sets can be combined after performing a within-set normalization. One simple normalization, the z -score, involves modifying the spike count values within each set (i.e., for each stimulus condition) by subtracting the mean and dividing by the SD for that set of responses. The subtraction eliminates the mean stimulus-evoked portion of the response, and the division scales the variance around the mean so that random fluctuations at high firing rates (which are known to be larger than those at low rates) are not unduly weighted. Further empirical justification for this normalization comes from the observation that r_{SC} changes very little with firing rate or stimulus condition, as shown in Results. The resulting z -scores can be represented in the order in which they occurred in the original experiment by the sequences z_1^i and z_2^i , $1 \leq i \leq M$, where M is the total number of trials in the experiment. Because $E z_1 = E z_2 = 0$, and $\sigma_{z_1} = \sigma_{z_2} = 1$, the equation for the correlation coefficient, Equation 3, simplifies to:

$$r_{\text{SC}} = E[z_1^i z_2^i]. \quad (4)$$

Table 1. Summary of notation regarding correlation

| Symbol | Description |
|--------------------------------|---|
| r | Pearson's correlation coefficient (Eq. 3) |
| r_{signal} | r for mean responses, computed from direction tuning curves |
| r_{noise} | r for trial-to-trial fluctuations around mean response |
| r_{SC} | r_{noise} computed directly from spike counts |
| $\text{TCC}(\phi)$ | Trial-based cross-covariance, generalizes r_{SC} beyond simultaneous trials (Eq. 5); $\text{TCC}(0) = r_{\text{SC}}$ |
| $\text{TCC}_{\text{hp}}(\phi)$ | TCC computed after z -scored spike counts are high-pass filtered |
| r_{ST} | Short-term component of r_{noise} , $\text{TCC}_{\text{hp}}(0)$ |
| r_{LT} | Long-term component of r_{noise} computed from smoothed TCC |
| $\text{TAC}(\phi)$ | Trial auto-covariance; $\text{TAC}(0) = 1$ |
| r_{AC} | Estimate of long-term correlation from smoothed $\text{TAC}(\phi)$ as $\phi \rightarrow 0$ |
| $C_{jk}(\tau)$ | Raw cross-correlation of simultaneous spike trains, averaged over trials (Eq. 14) |
| $\text{CCG}(\tau)$ | Spike train cross-correlogram, based on C_{jk} (Eq. 6) |
| $\text{ACG}(\tau)$ | Spike train auto-correlogram, based on C_{jk} , $j = k$ |
| $r_{\text{CCG}}(\tau)$ | r_{noise} estimated from the ACGs and CCG (Appendix A) |
| $S_{jk}(\tau)$ | Raw cross-correlation of post-stimulus time histograms, related to shift-predictor (Eq. 15) |
| $C_{jk}^*(\tau)$ | Raw all-way shift-predictor (Eq. 8) |

For a single set of paired responses, this equation is equivalent to Equation 3 because neither subtraction nor division by a positive constant (applied to the spike count data) changes the value of the correlation coefficient. For multiple sets of responses, the equation provides an aggregate correlation coefficient. Equation 4 can be generalized from responses that occurred on the same trial to responses that occurred on trials separated in time by a lag, ϕ , in units of experimental trials (~ 5 sec per unit; see below). This generalization, which we will refer to as the trial cross-covariance (TCC), is simply the cross-correlation of z_1 and z_2 :

$$\text{TCC}(\phi) = E[z_1^i z_2^{i+\phi}]. \quad (5)$$

The value at $\phi = 0$ is equal to r_{SC} (Eq. 3) averaged across all stimulus conditions (with appropriate weighting for the number of trials for each condition), and we use the symbols $\text{TCC}(0)$ and r_{SC} interchangeably. For $\phi \neq 0$, $\text{TCC}(\phi)$ is the correlation coefficient, with values from -1 to 1 , for temporal offsets in arbitrary numbers of trials. For a pair of uncorrelated neuronal responses, $\text{TCC}(\phi)$ will approach zero everywhere as the number of trials used in its estimate increases. The trial auto-covariance, $\text{TAC}(\phi)$, is defined in a similar manner by replacing z_2 with z_1 (or vice versa) in Equation 5, and by definition it is equal to unity for $\phi = 0$.

We already know that TCCs will have positive values at $\phi = 0$ for neuronal pairs with $r_{\text{SC}} > 0$. If interneuronal correlation arises at a time scale shorter than the trial duration, the positive value at $\phi = 0$ will stand as a narrow, isolated peak. However, if the correlation between neurons arises from slow changes in their responsiveness, the positive value at $\phi = 0$ will be part of a broader peak, i.e., the TCC will have positive values for $\phi \neq 0$ as well. Similarly, the presence of a broad peak around the origin in the TAC will indicate the presence of slow variations in the excitation of individual neurons.

Our use of the TCC does not rest on whether the horizontal axis is given in units of time or trials. We retained “trials” as the axis unit to avoid the technical difficulty associated with the cross-correlation of data sampled at somewhat irregular time intervals. The irregularity in the mapping from trials to time was caused by the monkey's failure to fixate immediately on 10–20% of trials during a recording session. To estimate the time scale of slow correlation, we will convert from trials to time using the average time between trial starts, ~ 5 sec.

Previous studies by Eggermont and Smith (1995, 1996) have attempted to separate correlation at multiple time scales using a method similar in concept to the TCC, but their time unit was 50 msec, roughly two orders

of magnitude faster than ours. Variations in firing rate on the time scale of 10s to 100s of milliseconds (Nelson et al., 1992; Eggermont and Smith, 1995; Arieli et al., 1996) are considered by us to be short term because they fall well within the duration of our behavioral epoch.

The spike train cross-correlogram. We measured correlation at the time scale of milliseconds using spike train auto- and cross-correlograms (ACGs and CCGs) (Perkel et al., 1967a,b). Our CCG is defined based on the trial-averaged cross-correlation, $C_{jk}(\tau)$ (defined in Appendix A, Eq. 14), of binary spike sequences from neurons j and k (typically, 1 and 2). In particular:

$$\text{CCG}(\tau) = \frac{C_{jk}(\tau)}{\Theta(\tau) \sqrt{\lambda_j \lambda_k}}, \quad (6)$$

where λ_j and λ_k are the mean firing rates (in spikes per second) of neurons j and k . For ACGs, $j = k = 1$ or 2 . The function $\Theta(\tau)$ is a triangle representing the extent of overlap of the spike trains as a function of the discrete time lag τ , i.e.:

$$\Theta(\tau) = T - |\tau|, \quad (-T < \tau < T), \quad (7)$$

where T is the duration of the spike train segments used to compute C_{jk} . Dividing C_{jk} by $\Theta(\tau)$ in Equation 6 changes the units of our CCG from raw coincidence count to coincidences per second and corrects for the triangular shape of C_{jk} caused by finite duration data.

In Equation 6, we chose to divide by the geometric mean spike rate (GMSR), $\sqrt{\lambda_j \lambda_k}$, because under this normalization the area of our CCG peaks remained relatively constant as firing rate varied [shown later; see also Krüger and Aiple (1988)] and because it is symmetric with respect to the two neurons. With this normalization, the CCG is the ratio of a coincidence rate to a mean spike rate and ends up with units of coincidences per spike. Once the shift-predictor (below) is subtracted, this normalization is similar to that of many other studies (Mastronarde, 1983a; Krüger and Aiple, 1988; Eggermont and Smith, 1996; Cardoso de Oliveira et al., 1997) and is conceptually similar to that proposed by Aertsen et al. (1989) for their “joint peri-stimulus time histogram.” A different normalization, dividing by the product of the spike rates, has been favored less often (Melssen and Epping, 1987, their Eq. 17; Das and Gilbert, 1995), and for our data was less appropriate than dividing by the GMSR.

Shift- (also known as shuffle-) predictors [defined in Perkel et al. (1967b)] for CCGs and ACGs were computed using the same normalization as above but based on the average cross-correlation of all $M^2 - M$ pairings of nonsimultaneous responses from neurons j and k for a set of M trials. This “all-way” cross-correlation, denoted $C_{jk}^*(\tau)$, can be computed efficiently from the cross-correlation of the post-stimulus time histograms (PSTHs), S_{jk} (defined in Appendix A, Eq. 15), according to the following expression:

$$C_{jk}^*(\tau) = \frac{MS_{jk}(\tau) - C_{jk}(\tau)}{M - 1}, \quad (8)$$

which approaches $S_{jk}(\tau)$ as M increases (Perkel et al., 1967b). Substituting C_{jk}^* for C_{jk} in Equation 6 gives the final shift-predictor. A shift-predictor computed from responses to ensemble stimuli (i.e., those that resulted from different sequences of random numbers; see above) will be referred to as an ensemble shift-predictor. When computed for replicate stimuli (i.e., repetitions of identical stimuli), it will be referred to plainly as a shift-predictor.

CCGs, ACGs, and shift-predictors were computed from data in the post-stimulus onset period 300–2000 msec to avoid processing the initial transient response. This made shift-predictors flatter and prevented changes in correlation strength that might be associated with the stimulus onset transient from influencing the analysis. We computed all quantitative results for the full trial as well and found only negligible differences. CCGs and ACGs were computed individually for each stimulus condition, shift-predictors were subtracted, and then averages were taken across all valid stimulus conditions. We set criteria for the minimum quantity of data required for neurons to be accepted into the CCG and ACG analysis pool. These rules were applied in order: (1) no trial was valid that had fewer than four spikes within the analysis window, (2) no stimulus condition was valid that had fewer than four valid trials or <64 spikes in total per neuron, and (3) no pair of neurons (or neuron) was included that had fewer than four valid stimulus conditions. These criteria eliminated 1 of 104 pairs from our direction tuning data set and 2 of 50 pairs from our coherence series data set.

Table 2. Summary of counts of cells and pairs

| Fixation task, direction tuning data | |
|--|--|
| 196 | Neurons with valid direction tuning data |
| 163 | Neurons classified as directional |
| 104 | Pairs with valid direction tuning data |
| 103 | Pairs met minimum data criteria for CCG computation |
| 69 | Pairs with both neurons directional |
| 57 | Directional pairs with $\Delta_{PD} < 90^\circ$ |
| Discrimination task, coherence series data | |
| 46 | Pairs with valid psychophysical discrimination data (83 neurons) |
| 72 | Neurons have valid neurometric thresholds, a (from 41 pairs) |
| 48 | Pairs met minimum data criteria for CCG computation (86 neurons) |
| 46 | Pairs with direction tuning data (85 neurons) |
| 34 | Pairs with both neurons directional |
| 29 | Directional pairs with $\Delta_{PD} < 90^\circ$ |

RESULTS

Our findings are organized as follows. The first section provides a brief description of our data for a typical pair of neurons and shows how all pairs are distributed according to the strength of their interneuronal correlation and the similarity of their directional tuning curves. The second major section is devoted to measuring the time scale of interneuronal correlation, which involves (1) separating long- and short-term correlation, (2) assessing the time scale of short-term correlation using spike train CCGs, and (3) relating CCG peaks to spike count correlation. A more efficient metric for spike count correlation is derived here and in Appendix A. The next major section of results reports the dependence of correlation, or synchrony, on stimulus parameters and on the decision-making and behavioral state of the animal. A brief section shows that neurons do not cluster with respect to their sensitivity to the stimulus or their relationship to behavior, and the final section describes control experiments for the influence of stimulus variance on our estimates of correlation.

Basic measurements of response correlation

Our results are based on simultaneous recordings from 107 pairs of MT neurons in three monkeys. We obtained directional tuning data for 104 pairs; we gathered discrimination data for a subset of 46 pairs. All recordings admitted to our database conformed to two requirements: both neurons were well isolated for at least 10 repetitions per stimulus condition, and at least one of the neurons yielded reliable, directionally selective responses to fully coherent random dot stimuli. For analyses involving CCG and ACG computations, we further restricted the database to pairs that satisfied criteria for a minimum number of spikes (see Materials and Methods). For ease of reference and consistency checking, the numbers of cells and pairs qualified for the major analyses are summarized in Table 2.

Figure 1 depicts a complete set of measurements for a representative pair of simultaneously recorded MT neurons. *A* and *B* are direction tuning curves for neurons 1 and 2, respectively. Both neurons were directionally selective and exhibited similar preferred directions and tuning bandwidths. *C* and *D* depict responses of the same two neurons as a function of motion coherence for both the preferred and null directions of motion. For these measurements, the preferred direction was set to 90° , approximating the optimal directions of both neurons. Off-line analysis of the data in *A* and *B* revealed the preferred directions

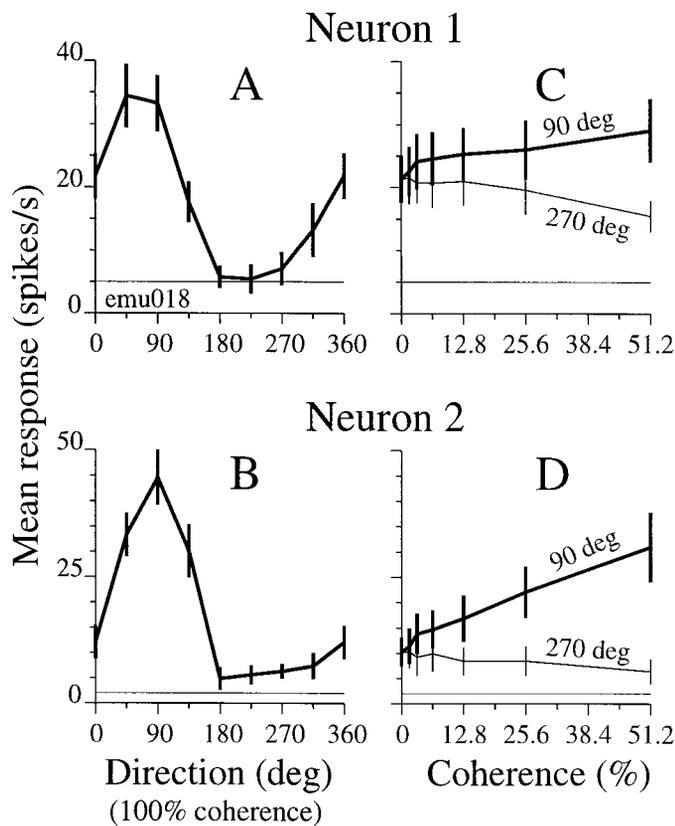


Figure 1. Example of direction tuning curves and responses versus coherence for a pair of neurons recorded simultaneously. *A* and *B* show mean firing rate as a function of stimulus direction for the two neurons (pair *emu018*). Stimuli were 100% coherence moving dots. Error bars show ± 1 SE. *Thin, flat lines* show spontaneous firing rate. *C* and *D* show mean firing rate as a function of coherence for preferred direction (90°, *thick lines*) and null direction (270°, *thin lines*) stimuli for the same neurons as *A* and *B*, respectively. Note that the minimum and maximum spike rates here do not reach those in *A* and *B* because 100% coherence was not included in these direction-discrimination experiments. Although some error bars are occluded at low coherence for null direction stimuli, they were roughly the same size as those for the preferred direction. A linear horizontal axis is maintained to emphasize the roughly linear relation between firing rate and coherence.

to be 58° and 82° for neurons 1 and 2, respectively (see Materials and Methods). In *C* and *D*, the firing rates of both neurons increased roughly linearly with motion coherence in the preferred direction and decreased linearly with motion coherence in the null direction, a typical pattern for MT neurons (Britten et al., 1993).

Using the direction tuning data for each pair of neurons, we assessed the strength of two distinct types of correlation, that of the mean responses and that of the variations about the mean. The former, commonly known as signal correlation, measures the similarity of tuning curves for a pair of neurons and was computed here as the correlation coefficient, r_{signal} , between the sets of data points from the direction tuning curves. For the curves in Figure 1, *A* and *B*, r_{signal} was 0.88, indicating a high degree of match. The distribution of r_{signal} for all pairs (Fig. 2*A*) was comparable to that of a more conventional but less general metric, the difference between preferred directions, Δ_{PD} , shown for comparison in Figure 2*B*. The dominant modes in both distributions, i.e., high r_{signal} and low Δ_{PD} , indicate that adjacent neurons in our study tended to have similar direction tuning, consistent

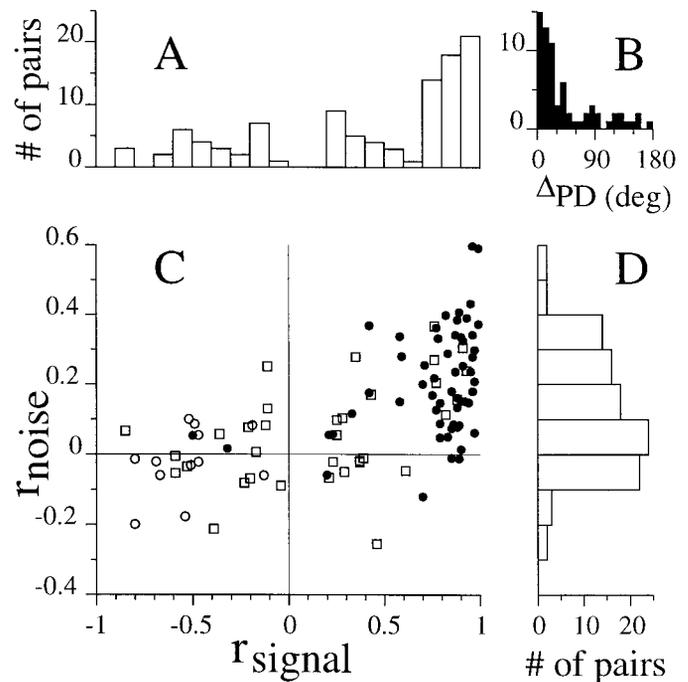


Figure 2. The joint distribution of signal and noise correlation for pairs of MT neurons. *A*, The marginal frequency distribution of r_{signal} for all pairs represented in *C*. *B*, The frequency distribution of Δ_{PD} (applicable to only those 69 pairs with both neurons directional) was similar to that reported by Albright et al. (1984) for pairs of MT neurons recorded successively at 50 μm intervals. Their study showed a second, small mode in the distribution for Δ_{PD} between 120 and 180°, corresponding to nearly opposite preferred directions for the two neurons. This mode was not readily apparent in our data; however, our sample size was smaller. From 104 pairs, 163 of 196 individual neurons (84%) were directional by the likelihood test described in Materials and Methods. *C*, Noise correlation is plotted against signal correlation for 103 pairs (direction tuning data). *Squares* indicate that at least one cell in the pair did not meet the directionality criterion ($n = 34$). *Circles* indicate directional pairs, and *filled circles* ($n = 57$) indicate pairs with similar preferred directions, i.e., $\Delta_{\text{PD}} < 90^\circ$. For the 57 directional pairs with similar preferred directions, r_{signal} was typically high (median 0.86), and the mean r_{noise} value was 0.20 (SD 0.15), which was significantly greater than zero ($p < 10^{-6}$; *t* test). The *two filled circles* in the *top left quadrant* represent pairs for which Δ_{PD} was only slightly $< 90^\circ$ (87 and 83°) and for which peculiarities of the direction tuning curves caused r_{signal} to be negative. For directional pairs with $\Delta_{\text{PD}} \geq 90^\circ$, $r_{\text{noise}} = -0.02$ (SD 0.09; $n = 12$; not different from zero; $p = 0.59$). For pairs with at least one nondirectional cell, $r_{\text{noise}} = 0.06$ (SD 0.15; $n = 34$; not different from zero; $p = 0.11$). *D*, Marginal frequency distribution of r_{noise} for all pairs represented in *C*.

with the known columnar organization of MT (Albright et al., 1984; DeAngelis and Newsome, 1999).

The second type of correlation is assessed not from the mean responses for all stimuli but from the trial-to-trial fluctuations (evidenced by the error bars in Fig. 1) around the mean response for each stimulus condition. This interneuronal correlation has therefore been dubbed noise correlation (Gawne and Richmond, 1993; Gawne et al., 1996; Lee et al., 1998). Noise correlation, or r_{noise} , is typically estimated by computing the correlation coefficient, r_{SC} , between the number of spikes generated by one of the neurons and the number of spikes generated by the second, simultaneously recorded neuron for a set of nominally identical stimuli. However, we developed a lower-variance estimator for r_{noise} (introduced and described in detail in the next section of Results) and have plotted those estimates against the values of r_{signal} for all pairs in Figure 2*C* (the marginal distribution of r_{noise}

is shown in *D*). The pairs appear to fall into two general groups in *C* that are not apparent from the marginal distributions alone. One group consists of pairs with very similar direction tuning curves (i.e., high r_{signal} values) and positive correlation. A second group consists of pairs with low or negative signal correlation and noise correlation near zero. Overall, the correlation coefficient between r_{signal} and r_{noise} is 0.61 ($p < 10^{-6}$; $n = 103$; direction tuning data).

The correlation of r_{noise} with r_{signal} is consistent with the notion that shared common input endows nearby neurons with similar tuning properties and makes them subject to similar noise sources. This observation is not unique to our data set, but it allowed us to focus our investigation of interneuronal correlation, when appropriate, on the cluster of neurons associated with non-zero r_{noise} values. We will use $\Delta_{\text{PD}} < 90^\circ$ as a criterion for making this separation.

The time scale of interneuronal correlation

In this section, we determine the time scale at which interneuronal correlation arises. We will quantify fluctuations in the neuronal response at time scales much slower and faster than the psychophysical trial and will show that the magnitude of r_{noise} for our MT pairs can be accounted for by the central peaks in their spike train CCGs on the order of 10s of milliseconds wide.

Short-term and long-term correlation

Since the earliest attempts to estimate r_{noise} in visual cortex, it has been recognized that slow processes could play an important role in determining its magnitude (van Kan et al., 1985; Bach and Krüger, 1986). Changes in neuronal excitation caused by motivational or attentional factors or fatigue could create a correlation at a time scale of anywhere between seconds and many minutes across a large population of neurons. On the other hand, common synaptic input to multiple neurons that operates on a millisecond time scale would also contribute to interneuronal correlation but across a smaller population of neurons sharing similar tuning properties. Because knowing the time scale of correlation may shed light on its origin and on its effect on pooled signals, our first goal was to determine to what extent long-term correlation was present and to calculate the remaining short-term component of r_{noise} once any long-term fluctuation of the firing rates was factored out. Assessing the presence of slow covariations in firing rate is also important because such covariation, when combined with faster stimulus-locked modulation, can lead to narrow CCG peaks that may be misinterpreted as evidence for fast synchronization (Brody, 1998, 1999). To tackle the problem of estimating slow changes in neuronal excitation for data collected in discrete epochs, i.e., trials, we developed a method called the TCC.

The TCC is a spike-count (as opposed to spike-train) -based cross-covariance that operates on the deviations from the expected responses (instead of the actual responses) for the two neurons, given the stimulus. Figure 3 outlines the TCC computation for two pairs of neurons, one for which correlation was predominantly long term, exceeding the duration of the 2 sec trial (*left column*), and a second for which correlation was predominantly short term (*right column*). The *top panels* (*A* and *D*) show for the individual neurons the z-score-normalized spike counts (see Materials and Methods) for trials in the order in which they occurred in the discrimination experiments. These traces estimate the levels of relative responsiveness of the neurons throughout the experiment. Beneath them, their auto-covariance functions, $\text{TAC}(\phi)$, are shown side-by-side (*B* and *E*; ϕ has units of

experimental trials, typically 5 sec per trial). Only the left or right halves of the TACs are shown (the functions are symmetrical about the origin), and the unity values at the origin are omitted. The gradual rise to a positive value around the origin, which was typical for our neurons, indicated that responses, or more precisely, response deviations from the mean, on any particular trial were correlated to those on earlier trials. For the two example pairs, the cross-covariance functions, $\text{TCC}(\phi)$, for the data in the *top panels* are shown at the *bottom* (*C* and *F*). $\text{TCC}(0)$ is the traditional interneuronal correlation coefficient for spike counts, r_{SC} (or more generally r_{noise}), whereas $\text{TCC}(\phi \neq 0)$ is a generalization of r_{SC} to responses occurring ϕ trials apart. In Figure 3*C* the broad central rise in the TCC indicates that the positive correlation on simultaneous trials ($\text{TCC}(0) = 0.1$; indicated by the circled dot) is related to a correlated drift in the activities of the neurons on a time scale longer than one trial. The example in *F* shows an entirely different outcome. Namely, the positive correlation coefficient for simultaneous trials does not extend to neighboring trials, despite the slow drifts in responsiveness of the two individual neurons evident from positive values near the origin in their TACs.

The TCC provides a framework for estimating long- and short-term components of r_{noise} , which is represented at $\text{TCC}(0)$. Long-term correlation, r_{LT} , is the value of the TCC around, but not at, zero. We estimated r_{LT} by replacing the value at zero with the average of its neighbors (at lag ± 1 trial), convolving with a Gaussian of SD four trials, and reading off the new value at zero (very similar results held for Gaussian SD two or eight trials). The traces from which r_{LT} was measured are shown as *smooth curves* superimposed on the raw TCCs (which still have their central values intact) in Figure 3, *C* and *F*. For the neuronal pair in *C*, r_{LT} was nearly the same as the raw r_{noise} value (the circled point is near the smooth line at lag 0), whereas in *F*, r_{LT} is close to zero and does not account for the value of r_{noise} . We used the same method (replacing the center and smoothing) to compute the long-term component of the auto-covariance, r_{AC} , from the two-sided, symmetrical forms of the TACs (Fig. 3*B*, *E*, arrows mark values).

To estimate the short-term component of r_{noise} , we removed the slow changes in responsiveness underlying r_{LT} by applying an ideal high-pass filter to the z-scored spike counts. The filter's cutoff frequency, 0.1 trial^{-1} (cutoff period 10 trials), was chosen to be faster than the mean time scale of slow changes in excitability observed in the TACs. The filtered data were subsequently renormalized to z-scores and used to compute a TCC (denoted TCC_{hp}) whose zero-lag value was our estimate of short-term correlation, i.e., $r_{\text{ST}} = \text{TCC}_{\text{hp}}(0)$. Figure 4 depicts the TCC and TCC_{hp} (*A* and *B*, respectively) for a pair of neurons that had substantial long- and short-term correlation. The long-term correlation was no longer visible in TCC_{hp} , but a narrow, central peak remained. A simpler approach to computing r_{ST} is to subtract r_{LT} from r_{SC} , i.e., from $\text{TCC}(0)$. However, this may yield less accurate results for many neurons because it is not in general correct to assume that r_{ST} and r_{LT} are additive.

Figure 4, *C* and *D*, shows database averages for our estimates of long- and short-term correlation. Separate averages are shown for pairs with $\Delta_{\text{PD}} < 90^\circ$ (*black bars*) and pairs with $\Delta_{\text{PD}} \geq 90^\circ$ or in which one neuron was not directional (*white bars*). A distinction between coherence series data (*C*) and direction tuning data (*D*) was maintained because we collected fewer total trials (typically 80) and had more pairs for direction tuning experiments ($n = 104$) than for discrimination experiments (at least 210 trials; $n = 48$). The database averages led to three significant observations.

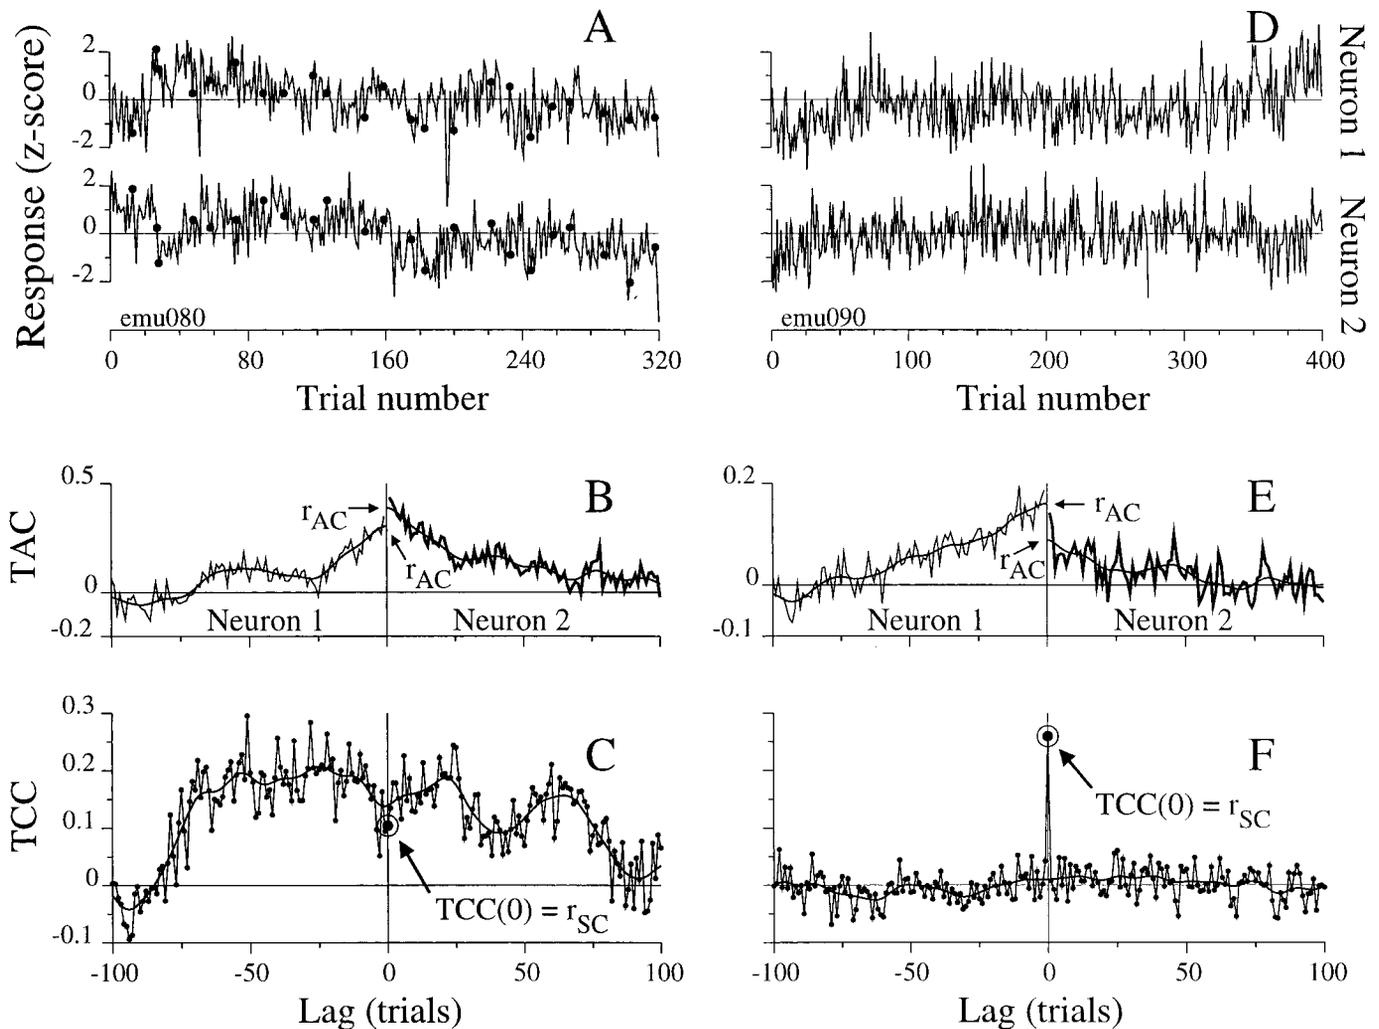


Figure 3. The trial cross-covariance (*TCC*) can reveal the presence or absence of long-term correlation. *A*, The z-score normalized spike counts for all 320 trials are plotted in the order they occurred for two simultaneously recorded neurons (*emu080*). Spikes were counted during the 2 sec stimulus, but trials occurred on average 5 sec apart, so 100 trials represent ~ 8.3 min. The *dots* show data points for preferred direction, 100% coherence stimuli to demonstrate how trials from one stimulus condition are interleaved among all others. *B*, Trial auto-covariance (*TAC*) plots are shown for the sequences from *A*, one on the *left* and one on the *right* side of lag 0 (the *TAC* is necessarily symmetrical about 0). Both have peaks of correlation around zero, indicating that the responsiveness of these neurons was not independent from trial to trial. The value at zero, 1 by definition, was omitted. The smooth line was used to estimate the value that the plot approached near the origin (*arrows*), referred to as r_{AC} or the long-term auto-correlation. *C*, The *TCC* is the cross-correlation of the sequences in *A*. $TCC(0)$ (circled point) is the aggregate r_{SC} for the pair. For this pair, the value at $TCC(0)$, 0.1, is associated with a peak that extends over lags of ± 70 trials. The *smooth line* shows the *TCC*, with center value replaced (see Results), convolved with a Gaussian (SD 4 trials) and was used to estimate the long-term cross-correlation, r_{LT} . *D*, Traces similar to those in *A*, but for a different pair of neurons (*emu090*). Data are shown for the first 400 of 1320 trials. *E*, *TACs* for the traces in *D*. Both neurons show positive correlation around zero. *F*, The isolated peak at $TCC(0)$ (circled point) for this pair indicates that correlation was predominantly short term, i.e., not associated with drifts in the responsiveness of the cells at time scales longer than the trial duration.

First, the average long-term auto-covariance, r_{AC} , was positive (*gray bars*; averaged across all individual cells), indicating that responses of single cells were correlated on a time scale longer than the single trial. For coherence series data, the mean long-term auto-covariance was 0.14 (SD 0.12; $n = 86$), only 4 of 86 cells had negative values, and the average *TAC* peak width at half-height was 48 trials (SD 51), corresponding to no less than 4 min. Second, however, the long-term cross-correlation, r_{LT} , was on average no different from zero (*t* test; $p = 0.39$; coherence series data). For the coherence series data, the distribution of r_{LT} was roughly Gaussian with mean 0.01 (SD 0.07; $n = 48$). Third, r_{ST} accounted for roughly the entire magnitude of r_{noise} for pairs in which both neurons were directional and had $\Delta_{PD} < 90$. For other

pairs, r_{ST} was not on average significantly different from zero, consistent with Figure 2*C*.

These results have the potentially counterintuitive implication that two neurons have responses that are correlated with their own responses on later trials and with each other's responses on simultaneous trials but not with each other's responses on later trials. In other words, long-term auto-correlation and short-term cross-correlation exist in the absence of long-term cross-correlation. This situation could arise if the sources of variance that caused the long-term auto-correlation in the responses of the individual neurons were independent from each other and from the source of variance that caused the short-term cross-correlation. That long- and short-term correlation arise from

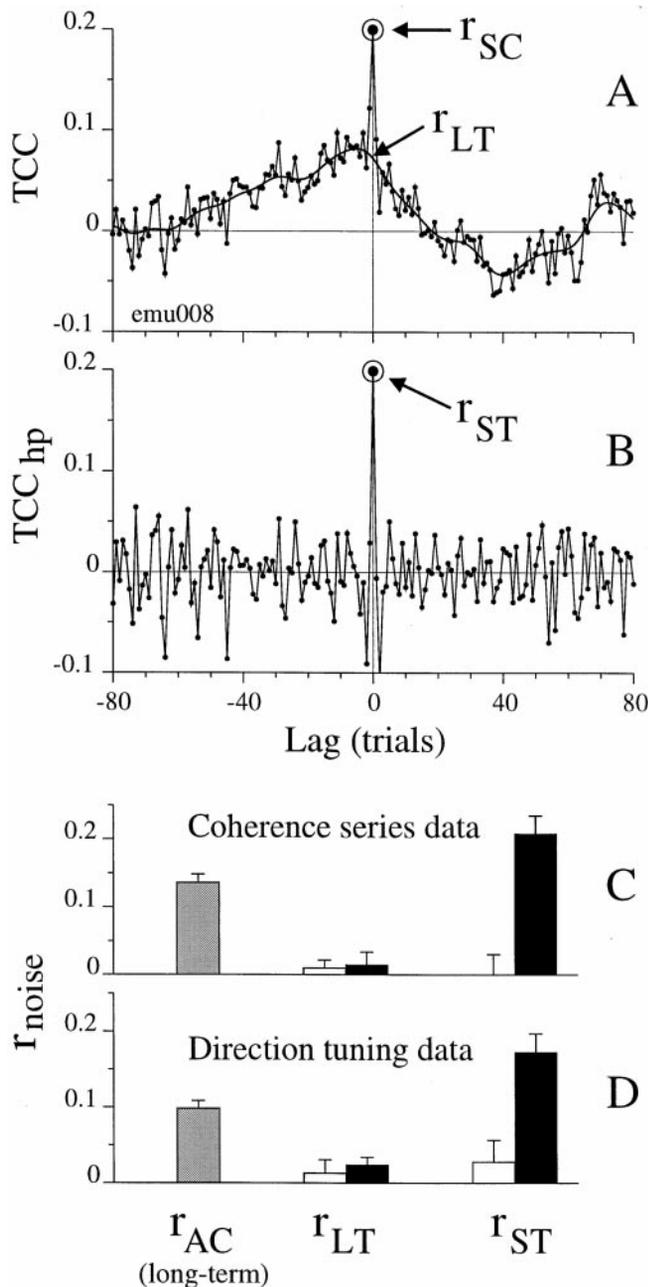


Figure 4. Noise correlation can be divided into short-term and long-term components. *A*, In the TCC (points connected by line segments, for pair *emu008*), the total noise correlation, r_{SC} , is the value at zero lag (circled point). The long-term correlation, r_{LT} , is the value of the smooth line (computed according to text, and Fig. 3C legend) at lag zero. *B*, For the same data as in *A*, the TCC_{hp} was computed after the z-scored spike counts were high-pass filtered. This eliminated the long-term correlation visible in *A*. The value at zero lag is taken to be the short-term correlation, r_{ST} . *C*, Database averages of correlation measures for coherence series data. Gray bar shows long-term auto-correlation, r_{AC} , averaged across all individual neurons ($n = 86$). Long- and short-term cross-correlation, r_{LT} and r_{ST} , are shown for directional pairs with $\Delta_{PD} < 90^\circ$ (black bars; $n = 29$) and other pairs (white bars; $n = 19$). Error bars show one SEM. *D*, The same measurements shown in *C* are shown here for direction tuning data from a larger set of pairs ($n = 196$ gray bars; $n = 57$ black bars; $n = 47$ white bars).

independent mechanisms would not be surprising, because they operate on time scales separated by four orders of magnitude, i.e., several minutes (shown above) versus 10s of milliseconds (shown in the next section).

In summary, slow drifts in the response strength of individual neurons were present ($r_{AC} > 0$) but on average uncorrelated ($r_{LT} \approx 0$) between pairs of neurons in MT, and therefore did not contribute significantly to the magnitude of interneuronal correlation across our database. Thus, r_{noise} was accounted for by the short-term component of correlation alone and must arise on a time scale no longer than the behavioral trial.

Spike train auto- and cross-correlograms

The positive value of r_{noise} (~ 0.21) associated with the cluster of points on the right side of Figure 2C did not result from long-term correlation, so we now test for its relationship to faster sources of correlation, the presence of which is revealed by spike train auto- and cross-correlograms. Examining ACGs as well as CCGs is important because ACGs bear on the interpretation of a CCG and because both are required for a mathematical result that we will use below to derive a new metric for r_{noise} . In this section, we establish, consistent with a body of cross-correlation studies, that correlation is largely limited in time to a small, central region of the ACGs and CCG and show that for our MT pairs there is a strong empirical relationship between that central region of the CCG and the traditional measure of spike count correlation, r_{SC} .

We computed the average spike train ACG for each individual neuron and the average CCG for each pair of neurons as described in Materials and Methods. Plots for one pair of neurons are shown on the left in Figure 5, and database summaries appear on the right. On the left, the ACGs (*A*) and CCG (*B*) are plotted in excess of the ensemble shift-predictor (see Materials and Methods) and are encased in lines showing ± 3 SD of the noise (estimated from the tails of the plots for lags from 400 to 800 msec). The ACG for neuron 1 (Fig. 5A, top trace, shifted vertically for visibility) has a dip near the origin, indicating that the likelihood of a spike occurring within 5 msec of another is lower than expected if spikes were fired independently of each other. This period of anti-correlation in the ACG is followed by a period of positive correlation from 7 until ~ 80 msec after a spike. Periods of both correlation and anti-correlation appeared in the ACG for neuron 2 as well (Fig. 5A, bottom trace). In addition, neuron 2 tended to fire pairs, or bursts, of spikes; however, this is not evident in the ACG plotted here because the positive values, at lags 2 and 3 msec, lay above the upper vertical limit of the plot and are not shown. The average CCG (Fig. 5B) for this pair of neurons had a central, somewhat asymmetric peak that did not extend beyond 100 msec from the origin.

Across our database, ACG shapes were diverse and varied in the presence and size of (1) a narrow central peak associated with short bursts, (2) a dip associated with a 1–3 msec absolute refractory period that was sometimes extended by a longer relative refractory or integration period (Abeles, 1982), and (3) a broader peak of positive correlation. The CCGs had mainly single, central peaks that varied in size, shape, and symmetry. Peak shapes were consistent with common synaptic input more so than with serial coupling (Moore et al., 1970). The shapes of our ACGs and CCGs were not consistent with the oscillatory Gabor functions that Kreiter and Singer (1996) used to describe CCGs in MT. In particular, we did not observe rounded central peaks flanked by similar but damped side-lobes.

We did not attempt a systematic classification of the subtleties of correlogram shapes, which would have required more data than we were able to collect for many of the pairs, but characterized only the extent in time of the correlation. This was accomplished for both correlation and anti-correlation by computing at

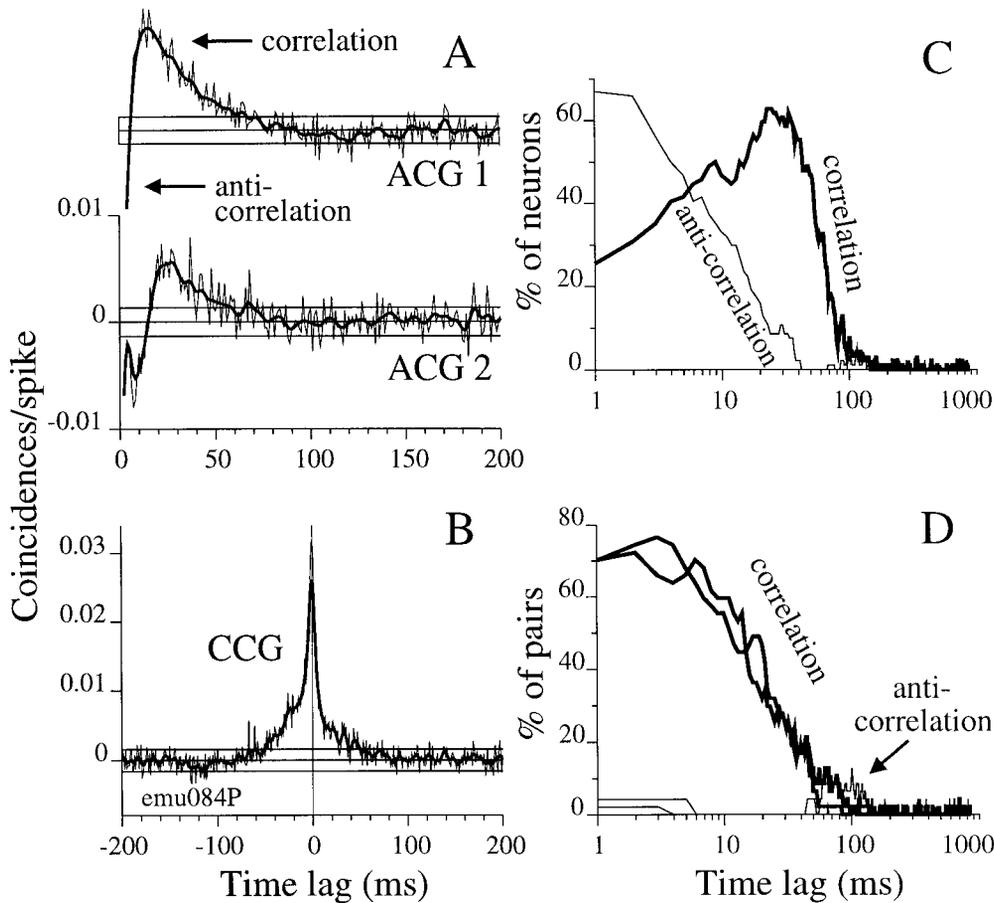


Figure 5. The time scale of spike train auto- and cross-correlation. *A*, ACG(τ) is plotted for two simultaneously recorded neurons (with ensemble shift-predictors subtracted). The *thick line* is a smoothed (Gaussian convolution, SD 2 msec) version of the raw ACG trace (*thin line*). The *three horizontal lines* show zero and ± 3 SD of the noise (see Results). The *top plot*, ACG 1, has been shifted upward for clarity. *B*, CCG(τ) is plotted for the pair with ACGs that are shown in *A*. Again, a smoothed trace is superimposed, and *horizontal lines* indicate zero and ± 3 SD of the noise. *C*, The fraction of neurons for which the ACGs were significantly (3 SDs) below (*thin line*) and above (*thick line*) the ensemble shift-predictor is plotted as a function of time lag (logarithmic axes). *D*, The fraction of pairs for which the CCG was significantly above (*thick lines*) and below (*thin lines*) the ensemble shift-predictor is plotted versus time lag. CCGs are two-sided, and results for negative time lags have been reflected onto positive time lags (yielding two traces at each thickness) for ease of comparison with the plot in *C*. Significant correlation occurs much more frequently for time intervals of several to tens of milliseconds than it does at or above 100 msec.

each millisecond time lag the fraction of cells that had correlation >3 SDs above the ensemble shift-predictor and the fraction that had anti-correlation <3 SDs below the shift-predictor. Correlograms were smoothed with a Gaussian of SD 2 msec before the test was applied to avoid counting isolated points that exceeded the criterion (as observed frequently in Fig. 5*A,B*). The results for the ACGs (Fig. 5*C*) revealed that significant response correlation for individual neurons was confined almost entirely to time lags <100 msec, was most prevalent around 30 msec, and decreased at shorter times because of the presence of anti-correlation associated with “non-burst” firing patterns or inter-burst intervals [described by Bair et al. (1994) for a comparable MT data set]. The extent of correlation in the CCGs is summarized in Figure 5*D* and, similar to that in the ACGs, was almost entirely confined to within 100 msec of the origin.

Two points deserve emphasis regarding these results. First, our analysis does not preclude weaker, yet significant, correlation that extends beyond 100 msec; it simply indicates that strong correlation, i.e., that which caused 3 SD differences between the correlograms and ensemble shift-predictors, was common at time scales on the order of 10s of milliseconds but was rare beyond 100 msec. Weaker, long-term sources of correlation certainly exist in MT but are not likely to contribute substantially to r_{noise} . Second, the time scale of correlation in our ACGs and CCGs is intrinsic to the visual system and does not result from temporal correlation in our stimulus because the signal strength (amount of preferred motion) in our dynamic dot stimulus was uncorrelated in time. In particular, the number of signal dots in any epoch (or in one video frame) was uncorrelated with that in any other epoch. The time scale of the correlation observed in Figure 5, *C* and *D*, matches

both the integration times for visual neurons upstream from MT (Hawken et al., 1996) and the temporal limits of motion perception for dynamic dot stimuli (Morgan and Ward, 1980). When analysis was restricted to zero coherence stimuli (which were effectively white noise to beyond 1 kHz), we found the same time scale of correlation across our database; therefore, the 45 msec time between signal dots in our stimulus was not responsible for the correlation observed here.

Having determined the typical time scale of correlation in our data, we may now apply a simple test to assess whether r_{noise} estimated in the traditional manner from the spike count for the entire trial is related to the central peak in the CCG. In Figure 6, the integral of CCG(τ) minus the ensemble shift-predictor (for $\tau = -32$ to 32 msec) is plotted against r_{ST} for our database (coherence series data). There is a clear relationship between these two measures of correlation (overall, $r = 0.76$, $p < 10^{-6}$, $n = 48$; for pairs with $\Delta_{\text{PD}} < 90^\circ$, $r = 0.71$, $p = 0.00001$, $n = 29$, *filled circles*; for other pairs, $r = 0.66$, $p = 0.002$, $n = 19$, *open circles*). This may seem striking because r_{ST} was derived from spike counts for the entire trial without information regarding the temporal structure of the spike trains, whereas the CCG area is based on the interrelationship of spikes occurring within 32 msec of each other. The significant positive correlation between the two metrics holds for limits of integration down to ± 2 msec ($r = 0.48$) but does not grow much in the range from ± 32 to ± 128 msec (e.g., $r = 0.80$ at both ± 64 and ± 128 msec). The data indicate that pairs of neurons with high spike count correlation also tend to have a substantial peak around the origin in their CCGs. This relationship is not given *a priori* (van Kan et al., 1985) and was not found in other studies of visual cortex (Gawne and

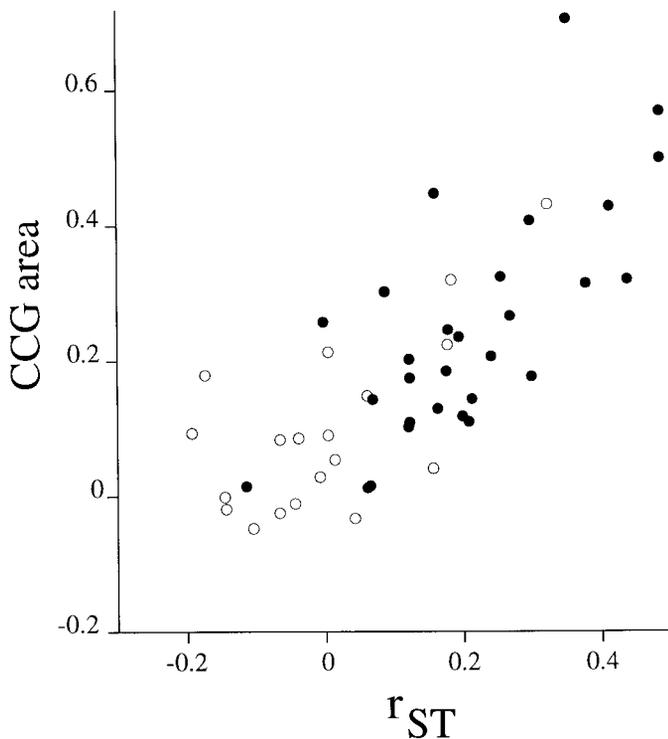


Figure 6. The area under the cross-correlogram peak is correlated with interneuronal correlation, r_{noise} . The area under CCG(τ) from $\tau = -32$ to 32 msec and in excess of the ensemble shift-predictor is plotted against r_{ST} , an estimator of r_{noise} . The correlation coefficient for all points is 0.76 ($p < 10^{-6}$; $n = 48$; coherence series data). A significant relationship also holds for the subset of 29 pairs for which $\Delta_{\text{PD}} < 90^\circ$ (●) and for the remaining 19 pairs (○; see Results for statistics). The ordinate and abscissa have similar values here, but their units are not the same. A metric that estimates r_{noise} by appropriate normalization of the CCG peak area is demonstrated in Figure 7.

Richmond, 1993; Gawne et al., 1996), although it was hinted at by Bach and Krüger (1986).

Assessing r_{noise} from the cross-correlogram

We will now make a more rigorous connection between r_{noise} and the area under the CCG by defining a metric based on the CCG that estimates exactly the value of r_{noise} under the condition that correlation has a limited time scale. Our approach derives from the fact that the equation for r_{SC} (the well known Pearson's correlation coefficient) can be rewritten in a form that is based solely on the areas under the spike train CCG and ACGs as follows (from Appendix A, Eq. 26):

$$r_{\text{SC}} = \frac{\text{Area under CCG}}{\sqrt{\text{Area under ACG}_1 \times \text{Area under ACG}_2}}, \quad (9)$$

where the areas are integrated across all lag times in the correlograms. However, if correlation is limited to short time lags, as suggested by results from the previous section, only those regions near the origin will contribute to non-zero areas in Equation 9. The flanks of the CCG and ACGs, which approach the shift-predictors, will contribute on average nothing but noise. We therefore propose the use of a metric, $r_{\text{CCG}}(\tau)$ (defined in Appendix A, Eq. 27), which estimates r_{noise} by integrating only a limited central region (from $-\tau$ to τ msec) of the CCG and ACGs. This metric eliminates the noise that would be contributed by the flanks of the correlograms by simply not including the flanks in

the integration. In essence, it assumes that the correlograms beyond $\pm\tau$ are on average equal to the shift-predictors.

Before applying the $r_{\text{CCG}}(\tau)$ metric to our MT data, we tested it on pairs of simulated spike trains that had a central, Gaussian-shaped CCG peak (SD 4 msec) and an r_{noise} value of exactly 0.2. For the simulated data, all of the area in the CCG (and ACGs) was concentrated near the center, and the expected value of the flanks (when the shift-predictor was subtracted) was known to be zero. Figure 7A shows $r_{\text{CCG}}(\tau)$ plotted for 10 sets of simulated spike trains (details of the simulation are given in the Figure legend). As τ increased, the average value of r_{CCG} increased until it reflected the true value, 0.2. A plateau occurred when τ exceeded the time scale of the correlation, and further increases in τ caused a loss of precision as noise from the ACG and CCG flanks was integrated. When τ reached the full trial duration (here 1700 msec), r_{CCG} became equivalent to r_{SC} , according to Equation 28. This simulation shows vividly how noise from the tails of the ACGs and CCG corrupts r_{SC} , and it demonstrates that a more accurate estimate of r_{noise} can be obtained with $r_{\text{CCG}}(\tau)$ when τ is shorter than the trial duration (but longer than the time scale of correlation).

We plotted $r_{\text{CCG}}(\tau)$ for our neuronal pairs and found a similar pattern of results. Curves for one pair are shown in Figure 7B for 11 coherence levels (from 100% preferred to 25.6% null direction motion, which satisfied the minimum data requirements stated in Materials and Methods). The curves increased together to $r \approx 0.16$ as τ approached 30–40 msec but then diverged as τ grew larger. This was consistent with the CCGs (data not shown), which had central peaks that fell to the level of the shift-predictor at ~ 30 –40 msec from the origin. The direction of divergence of curves such as those in Figure 7B typically did not depend on the stimulus condition (a systematic analysis is given in the next section), so we averaged across conditions to get an $r_{\text{CCG}}(\tau)$ curve for each pair, and we averaged across pairs to get one database curve. The database curve for pairs having $\Delta_{\text{PD}} < 90^\circ$ (Fig. 7C, filled circles) approached an asymptote of ~ 0.21 for values of τ above 32–64 msec. The value 0.21 was the same as that for short-term correlation for this database (Fig. 4C, right-hand bar), and the timing of the approach to the asymptote was consistent with the time scale of correlation observed in the ACGs and CCGs (Fig. 5C,D). Figure 7C also shows the SD for the r_{CCG} estimate (open circles, averaged across the same set of curves used to compute the mean). The SD grew with increasing τ even after the mean of $r_{\text{CCG}}(\tau)$ had leveled off. This shows the inefficiency of a long integration time such as that associated with the r_{SC} metric (i.e., the entire trial duration). Finally, a direct comparison of r_{SC} with $r_{\text{CCG}}(32)$ for individual pairs is provided in Figure 7D. The SD was always smaller for r_{CCG} (thick lines) than for r_{SC} (thin lines). Two points are labeled, one for the pair from B (*emu005*) and another (*emu080*) from Figure 3C, that had a large long-term component of r_{SC} . For the latter, $r_{\text{CCG}}(32)$ is much less than r_{SC} because $r_{\text{CCG}}(32)$ discounts long-term correlation. It does so by integrating area over only 1.9% (32 msec/1700 msec) of the CCG and therefore captures only 1.9% of the excess area that a source of long-term correlation spreads evenly across a CCG.

Clearly, r_{CCG} provided a more repeatable (less noisy) estimate of interneuronal correlation (for $\tau < T$) than did r_{SC} , but we wanted to verify that it also maintained the relationships that r_{SC} had with the measures for similarity of neuronal tuning mentioned above, namely, r_{signal} and Δ_{PD} . Compared with r_{SC} , $r_{\text{CCG}}(\tau)$ was more positively correlated with r_{signal} (Pearson's $r = 0.59$, rather than 0.53, for both $\tau = 32$ and 64 msec; $n = 46$) and was

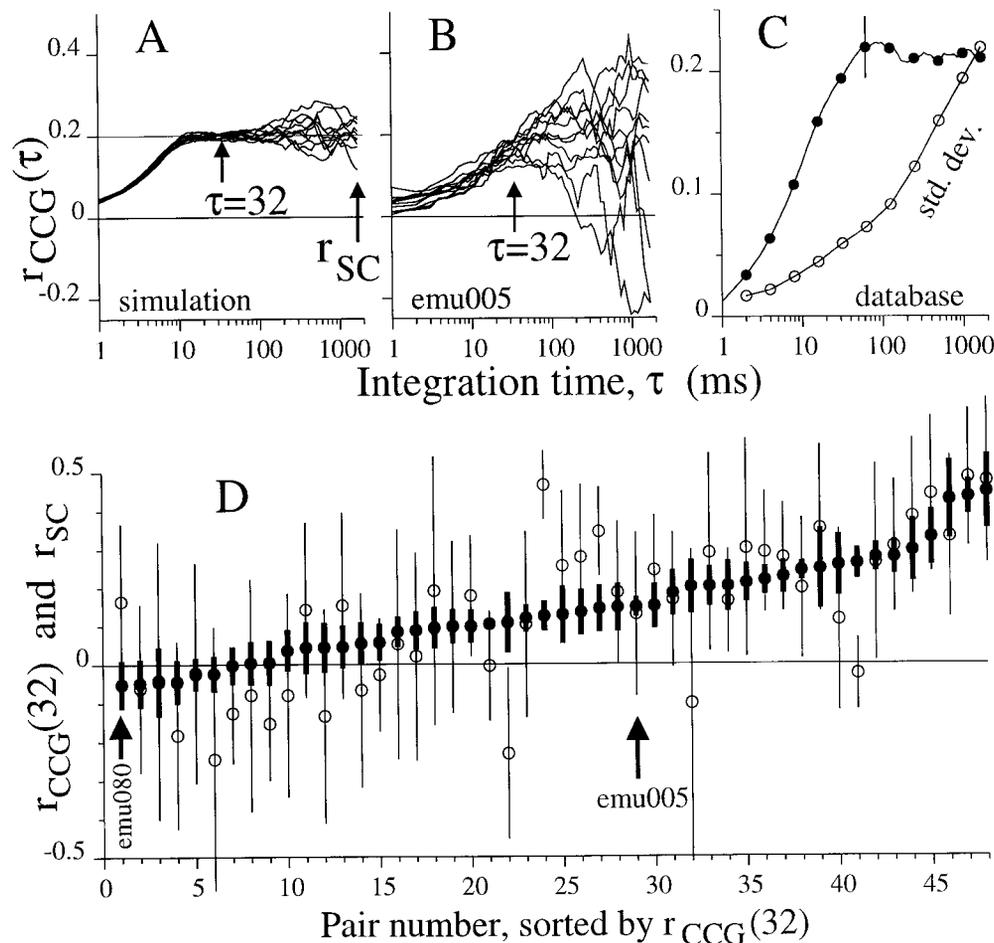


Figure 7. Computing $r_{CCG}(\tau)$ for simulated and neuronal data. *A*, $r_{CCG}(\tau)$ is plotted against integration time τ for 10 blocks of simulated pairs of spike trains in which the time scale and strength of correlation were known. (For each trial, 2 simulated spike trains were generated by selecting spike times independently and at random with probability 0.2 per spike from a Poisson spike train having mean rate 200 spikes per second. The rate is arbitrary and does not affect the results. The spike times in one of the resulting trains were jittered by adding a Gaussian random variable with mean zero and SD 4 msec. The resulting pairs of spike trains have $r_{noise} = 0.2$ and a time scale of correlation matching the Gaussian jitter.) For the simulated data, $r_{CCG}(\tau)$ approached the true value of r_{noise} once τ exceeded the time scale of correlation but became corrupted by noise as τ increased. Arrows indicate where $r_{CCG}(32)$ and r_{SC} are read out. Here, $r_{CCG}(32) = 0.200$ SD 0.009, whereas $r_{SC} = 0.197$ SD 0.037 ($n = 20$). The ratio of SDs was 1.4. *B*, Results from an analysis similar to *A* but for 11 coherence levels for neuronal data (pair *emu005*). As in *A*, the curves rose from zero as τ increased but became highly corrupted by noise as τ exceeded 100 msec. For this pair, the width at one-half height of the CCG (data not shown) was 9 msec. *C*, The mean value of $r_{CCG}(\tau)$ across 29 pairs (averaged across all coherence levels for each pair with $\Delta_{PD} < 90^\circ$) is plotted as a function of τ in 1 msec steps (filled circles mark octave steps). Vertical bar indicates ± 1 SE

at $\tau = 64$ msec. The SD of $r_{CCG}(\tau)$ is plotted at octave steps only (open circles connected by straight lines). The value of r_{CCG} reached an asymptote of ~ 0.21 for τ around 30–100 msec, but the SD continued to increase with τ . *D*, $r_{CCG}(32)$ (thick lines, filled circles) and r_{SC} (thin lines, open circles) for 48 pairs of neurons (coherence series data). Arrows indicate points corresponding to example pairs from *B* (*emu005*) and from Figure 3C (*emu080*). Points are sorted by increasing value of r_{CCG} . Vertical lines show ± 1 SD (computed across coherence levels) and were always larger for r_{SC} .

more negatively correlated with the logarithm of Δ_{PD} (Pearson's $r = -0.47$, rather than -0.36 , for both $\tau = 32$ and 64 msec; $n = 34$, where the logarithm was taken to correct the skew of the distribution in Fig. 2*B*).

In summary, it appears that $r_{CCG}(\tau)$ accurately captures the amount of interneuronal correlation for our pairs. That it does so for τ as small as 32 msec shows that most of the correlation observed at the time scale of the behavioral epoch can be accounted for by CCG peaks at a time scale nearly two orders of magnitude shorter. Therefore, mechanisms underlying narrow, central CCG peaks affect response properties relevant to both temporal and rate coding.

Dependence of correlation on stimulus and behavior

Assessing the dependence of correlation on stimulus parameters is necessary to justify averaging r_{noise} values and CCGs across stimulus conditions as we have done. In addition, this assessment is important with respect to both stimulus and behavioral parameters because of the potential link between correlation, or synchrony, and the perception of the animal as reflected by its behavior. Here we examine how correlation changes with the firing rates of the neurons, the direction and coherence of stimulus motion, and the presence of the stimulus, and we test whether synchronous activity exerts extra influence on the mon-

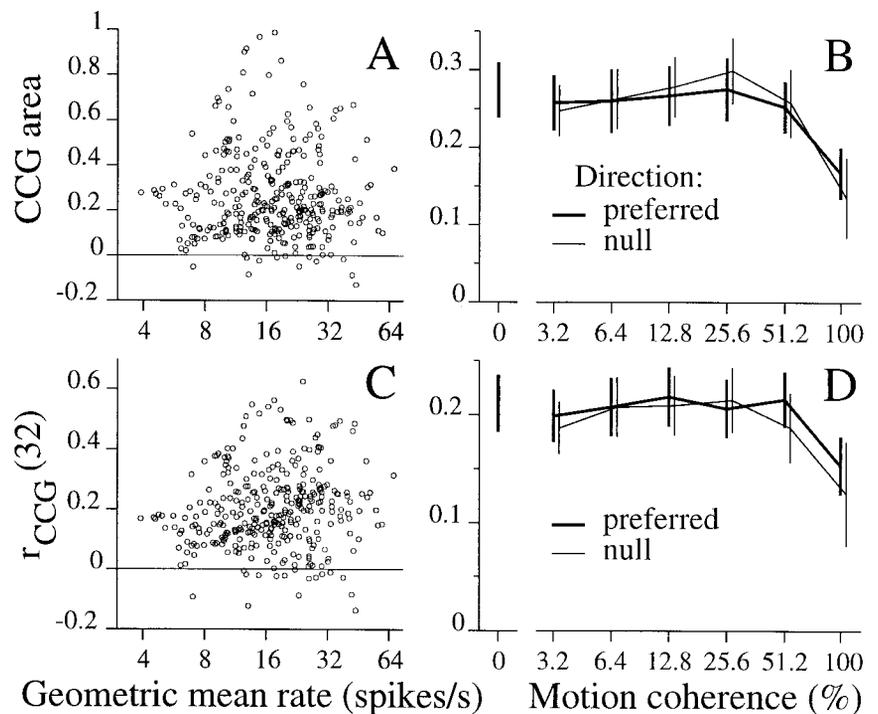
key's decision and whether it varies from passive fixation to active discrimination.

Correlation versus firing rate, direction, and motion coherence

Because firing rate varied as our stimulus parameters changed, we first established that our correlation metrics did not show a substantial dependence on firing rate before testing for more interesting relationships between interneuronal correlation and other variables. Figure 8, *A* and *C*, shows scatter plots of the area under the CCG peak (from -32 to 32 msec) and $r_{CCG}(32)$ versus geometric mean spike rate for each coherence level for the 29 directional pairs with $\Delta_{PD} < 90^\circ$. Firing rate was not significantly correlated with CCG area and showed only a weak relationship with r_{CCG} (see Figure legend for details). A pair-by-pair analysis also revealed no overall trend, although several individual pairs showed significant relationships (see Fig. 8 legend). Similar results held for data from the direction tuning experiments, for integration times ranging from several to hundreds of milliseconds, and when all pairs were included in the analysis.

The same two correlation metrics were largely constant when plotted against stimulus direction and coherence, except at 100% coherence where both measures were lower (*B* and *D* show CCG area and r_{CCG} , respectively, Fig. 8). The numbers of individual pairs for which these metrics were significantly correlated with

Figure 8. CCG peak area and r_{noise} plotted as a function of geometric mean spike rate (GMSR) and motion coherence. **A**, The area of the CCG (minus the ensemble shift-predictor) integrated from $\tau = -32$ to 32 msec is plotted as a function of GMSR for pairs of directional neurons with $\Delta_{\text{PD}} < 90^\circ$. Each point represents data for one coherence level from 1 of 29 qualified pairs. Pearson's correlation coefficient for this scatter was not significantly different from zero ($r = -0.02$; $p = 0.65$; $n = 329$). We chose our CCG normalization (Eq. 6) to realize this empirical observation. For individual pairs, 3 of 15 negative relationships and 6 of 14 positive relationships were significant ($p < 0.05$). **B**, Mean CCG area (± 32 msec) averaged over the same 29 pairs as in **A** is plotted for preferred (*thick line*) and null (*thin line*) motion as a function of coherence (there are typically 24–29 pairs per point because some pairs were not tested at all coherence levels). *Vertical bars* show ± 1 SEM. Values remained relatively constant except for a 43% reduction at 100% coherence compared with the average across all lower coherence stimuli (preferred and null directions combined). **C**, A measure of r_{noise} , here $r_{\text{CCG}}(32)$, is plotted against GMSR for the same set of pairs. There is a small overall positive correlation with mean spike rate ($r = 0.14$; $p = 0.01$; $n = 329$). For individual pairs, 1 of 12 negative and 6 of 17 positive relationships were significant. **D**, Like CCG area, the r_{noise} measure varies little with motion direction and coherence except at 100% coherence where it dropped by 32% relative to the average across all lower coherence levels (preferred and null directions combined). As in **B**, typically 24–29 pairs contributed data to each point.



coherence were almost identical to those for spike rate. The drop in correlation strength at 100% coherence can be related to the nature of MT responses to coherent and incoherent motion. MT neurons typically show clear stimulus-locked modulation for stimuli of $< 100\%$ coherence, but at 100% coherence there is little or no such modulation (Bair and Koch, 1996). How this modulation impacts our measures of correlation is the subject of the last section of Results. Whether the reduction in r_{noise} at 100% coherence is also related to a previous report that correlation is almost completely abolished during high contrast motion in MT (Cardoso de Oliveira et al., 1997) is discussed in the next section.

The consistency of r_{noise} in the face of large changes in firing rate indicates that the underlying mechanism did not act additively to alter neuronal firing rates, for if it did, r_{noise} would be larger at lower firing rates. In the absence of substantial overall relationships between our correlation metrics and the stimulus direction and coherence or the firing rate, we chose to average these metrics across all stimulus conditions. The observed decrease at 100% coherence had little influence on our statistics because $< 10\%$ of our coherence series data was collected at $c = 100\%$.

Correlation during spontaneous and stimulus-driven activity

We tested the dependence of correlation on the presence of the stimulus by computing r_{noise} for a 330 msec epoch of spontaneous activity and for an equal duration epoch of stimulus-driven activity. The spontaneous epoch began when the monkey acquired fixation and ended 30 msec after stimulus onset, precluding the arrival of stimulus-driven activity in MT (Raiguel et al., 1999). The driven epoch began 30 msec after stimulus onset. We limited analysis to pairs that had at least four stimulus conditions each having at least 10 trials with at least one spike per trial per cell during the 330 msec period. The value of $r_{\text{CCG}}(32)$ for the spontaneous epoch was significantly correlated with that for the

driven epoch ($r = 0.63$; $p = 0.00001$; $n = 40$), and the average difference between the values for spontaneous and driven activity, 0.018 (SD 0.14), was not significantly different from zero. Limiting the analysis to directional pairs with $\Delta_{\text{PD}} < 90^\circ$ gave nearly identical results. Similar results were found when (1) r_{ST} , computed from the TCC, was substituted for $r_{\text{CCG}}(32)$, or (2) the driven epoch was defined to be the entire stimulus epoch, rather than just the first 330 msec. We conclude that noise correlation during spontaneous activity is similar to and predictive of the noise correlation during activity evoked by our random dot stimuli.

This result stands in striking contrast to the report of Cardoso de Oliveira et al. (1997) that interneuronal correlation in MT is present during spontaneous activity but is practically abolished during visual stimulation. To determine whether our correlation values were more similar to their values for spontaneous or for driven activity, we normalized our CCGs according to their methods (after dividing by the geometric mean spike rate, we used a three-point boxcar function to smooth the CCGs and then found the peak within ± 100 msec of zero) and computed peak height, position, and width statistics like those presented in their Figure 5. All three measures from our data were well matched to their results for spontaneous activity, indicating that our results differ only during visual stimulation. If we assume that high-contrast, coherently moving stimuli reduce correlation strength between responses of nearby MT neurons, then it remains to be determined why our strongest stimulus (100% coherence motion) caused a decrease in correlation strength that was small compared with the decrease caused by the square-wave grating of Cardoso de Oliveira et al. (1997).

Does correlation change with behavior?

Investigators have hypothesized that synchronous firing among cortical neurons underlies various coding or processing functions

(for review, see Singer and Gray, 1995; Roelfsema, 1998). Our data provide the opportunity to determine whether synchrony among adjacent MT neurons is correlated with perceptual choice or behavioral state.

The relation of synchrony to perceptual choice is best assessed at low motion coherence where the monkey correctly identifies the direction of motion on some trials but makes mistakes on others. The psychometric function in Figure 9*A* (thick line, filled circles) plots the monkey's performance on trials for which the stimulus was optimized for the pair of neurons the tuning curves of which are shown in Figure 1. We asked whether synchrony was stronger on trials in which the animal chose the direction preferred by the pair of neurons, as might be expected if synchronously active neurons exert stronger effects on downstream decision circuitry. To test this, we divided the trials for each stimulus condition (i.e., for a particular coherence level and direction) into two groups, one in which the animal chose the preferred direction (for the pair) and one in which the animal chose the null direction. Note that one group corresponds to correct decisions, whereas the other corresponds to incorrect decisions (where the correspondence depends on whether the direction of motion was null or preferred for the stimulus condition) except at zero coherence where there was no "correct" response.

We considered only stimulus conditions that had at least 10 trials with preferred responses and 10 trials with null responses; therefore, 51.2 and 100% coherence conditions were rarely included because the monkey rarely made 10 mistakes for such salient stimuli. This limited the number of pairs for this analysis from 46 to 35. Figure 9*B* shows CCGs for preferred and null decision trials for the same pair of neurons illustrated in Figures 1 and 9*A*. The CCGs appear virtually identical, which was typical for our data set. Figure 9, *C* and *D*, depicts quantitative measurements of the area under the CCG from -32 to $+32$ msec (*C*) and from -2 to $+2$ msec (*D*) for preferred and null decision trials for 137 stimulus conditions from the 35 pairs of neurons. In both panels, the points cluster around the unity diagonals, showing that synchronous firing did not differ between the two decision states (paired *t* test; $p = 0.75$ for *C*, $p = 0.94$ for *D*). This result also held for the r_{CCG} metric, for all integration times tested (from ± 2 to ± 128 msec), and when only directional pairs were tested.

We also analyzed synchrony simply as a function of motion coherence, regardless of perceptual choice. At low coherence, the dot patterns appear to be a white noise stimulus and elicit no global motion percept. As coherence increases, however, observers perceive the entire stimulus to drift in the specified direction as though the disparate motion signals provided by individual dot pairs are bound into a perceptually coherent whole. Theories of perceptual binding that postulate a unique role for synchronous neural activity might predict that synchrony should be stronger for coherent ($c = 100\%$) than for incoherent dot patterns ($c = 0\%$). However, we have already seen that the opposite is true (Fig. 8*B,D*).

Finally, we compared CCGs obtained during passive fixation (direction tuning experiments) with those obtained during active discrimination to determine whether the overall behavioral state of the animal was correlated with neural synchrony. In the subset of experiments in which both blocks of data were obtained, the area under the CCG did not differ systematically between the two states (paired *t* test; $t = -0.06$; $p = 0.95$; $n = 46$), and the measurements were highly correlated between the two states ($r = 0.90$; $p < 10^{-6}$). In short, we found no evidence that synchronous

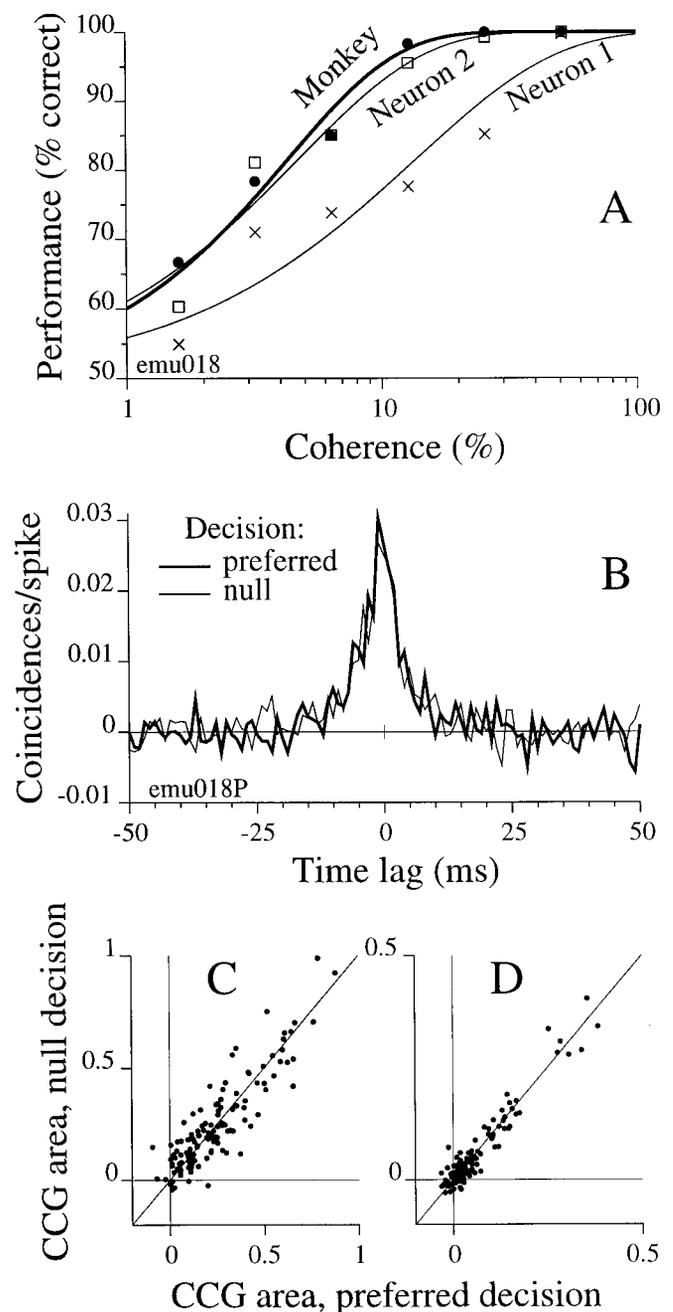


Figure 9. Psychophysical and neuronal performance. *A*, The monkey's performance (filled circles) and the neuronal performance for neuron 1 (\times 's) and neuron 2 (squares) are plotted as a function of motion coherence (logarithmic axis) for the same pair of neurons as in Figure 1. The lines show fits to Equation 2. The thick line is for the monkey's psychophysical responses. *B*, An example of average CCGs for preferred (thick lines) and null (thin lines) direction decisions. CCGs for coherence levels of ± 6.4 , ± 3.2 , and 0% were averaged together; other coherence levels did not have a sufficient number of choices in each direction. *C*, The area under the CCG between -32 and 32 msec and in excess of the shift-predictor for null decisions is plotted versus that for preferred decisions. Each point ($n = 137$) shows data for a particular coherence level and direction, so there are multiple points per neuronal pair ($n = 35$). *D*, For the same data set as in *C*, the comparison of CCG area is made for a narrower integration region, from -2 to 2 msec. Results in *B–D* reflect a lack of correlation between perceived direction of motion and the magnitude of synchronous activity in the population of neurons that prefer the perceived direction.

firing varied systematically as a function of perceptual decision or behavioral state.

Do sensitive or informative neurons cluster?

For each experiment in which we obtained psychophysical data, we used analytic methods based on signal detection theory [see Materials and Methods, or see Britten et al. (1992) for detailed methods] to compare the directional sensitivity of each neuron with the monkey's psychophysical sensitivity. Figure 9A illustrates the outcome of this analysis for the data depicted earlier in Figure 1, C and D. The *filled circles* represent the psychophysical performance of the monkey on the direction discrimination task, which increased from nearly chance at low coherence levels to perfection at the three highest levels. Psychophysical threshold, defined as the motion coherence that supported 82% correct performance, was 4.3% coherence. The *x*'s and *squares* indicate the performance of the two MT neurons measured on the same trials represented in the psychometric curve. Neuron 2 was as sensitive to the directional signals as was the monkey psychophysically, yielding a neurometric threshold of 4.7% coherence. Close correspondence between neuronal and psychophysical thresholds is common in MT (Britten et al., 1992). In contrast, neuron 1 was considerably less sensitive to motion signals in the displays, yielding a threshold of 13.8% coherence (sensitivity = 1/threshold). Across 72 directional neurons studied in 41 discrimination experiments, the geometric mean ratio of neuronal to psychophysical threshold was 1.72 (range, 0.27–11.5), a value higher than those previously observed in this laboratory (Newsome et al., 1989; Britten et al., 1992; Celebrini and Newsome, 1995). The discrepancy arises because the inclusion criterion for direction selectivity was less stern in the current study to maximize the number of pairs. Interestingly, neither neuronal thresholds nor choice probabilities [defined in Britten et al. (1996)] were significantly correlated between adjacent MT neurons in our sample. Thus we find no evidence for clustering of neurons that are particularly sensitive to the stimulus or that have particularly close relationships to behavior.

Controls for stimulus variance–replicate stimuli

Our estimates of interneuronal correlation have been based on responses to ensembles of stochastic stimuli in which the random detail of the dot patterns differed from repeat to repeat within a particular stimulus category. In principle, such sets of nonidentical stimuli could inflate r_{noise} estimates and increase CCG peak sizes if the responses of the neurons were influenced by the random variation across stimuli. For example, if 15 of 30 stimuli that were generated at 6.4% coherence had by chance slightly more motion in the preferred direction than the other 15 stimuli, an ideal pair of neurons with no common noise source but having identical direction preferences would tend to fire on average more for the former than for the latter 15 stimuli. This would yield an erroneous positive value of r_{noise} , which should otherwise be zero. Below we describe direct experimental controls as well as simulations that allow us to estimate the magnitude of this effect in our data.

In our experiments, random variation from stimulus to stimulus was necessary to prevent the monkeys from associating particular spatial patterns with a reward. However, for four pairs of neurons we interleaved experiments using replicate stimuli in which the dot patterns for a particular stimulus condition were identical (see Materials and Methods). Estimates of r_{noise} for the four controls using both the r_{SC} and $r_{\text{CCG}}(32)$ metrics are presented in Figure

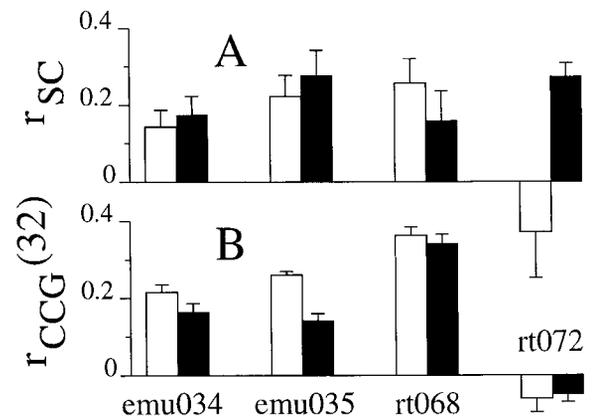


Figure 10. A comparison of interneuronal correlation strength for stochastic versus replicate stimuli. *A*, r_{SC} is plotted for four pairs tested with both ensemble (*white bars*) and replicate (*black bars*) stimuli. No significant change was observed for the first three pairs. For the fourth pair, the responses to replicate stimuli had a large long-term component of correlation because of drifts in firing rate during the experiment ($r_{\text{LT}} = 0.27$, $r_{\text{ST}} = 0.02$). Error bars show SE. *B*, $r_{\text{CCG}}(32)$ is shown for the same four pairs. In the first three cases, $r_{\text{CCG}}(32)$ was less for replicate stimuli. The difference was significant for *emu035* (see Results). Lower values for replicate stimuli (*black bars*) were consistent with the reduction in the CCG peak that occurred when the true shift-predictor (Fig. 11C) was subtracted. Cases *emu034* and *emu035* were based on coherence series data; *rt068* and *rt072* were based on direction tuning data ($c = 100\%$). For *rt072*, the low r_{noise} value was consistent with $\Delta_{\text{PD}} > 90^\circ$.

10. The values of r_{SC} (*A*) offered no evidence that interneuronal correlation was greater for stochastic stimuli (*white bars*) than for replicate stimuli (*black bars*), but the lower-variance estimates provided by $r_{\text{CCG}}(32)$ (*B*) painted a clearer picture. For pairs *emu034* and *emu035*, $r_{\text{CCG}}(32)$ was higher for stochastic stimuli than for replicate stimuli ($p = 0.08$ and $p = 0.00001$ respectively, *t* tests). For the other two control pairs, the difference was negligible.

An examination of the PSTHs, CCGs, and shift-predictors for *emu035* (the pair that had a significant decline in r_{CCG} for replicate stimuli) reveals how stimulus-locked modulation can inflate r_{noise} . For replicate stimuli, the stimulus-locked modulation of firing rate is captured in the PSTHs (Fig. 11*A,B*, *thin lines*, neurons 1 and 2, respectively), but when the stimulus varies from one repeat to the next (ensemble stimuli), the modulation is washed out (*A, B*, *thick lines*). The difference in the PSTHs carries over into the CCG shift-predictors because shift-predictors are closely related to the cross-correlation of the PSTHs [see Eq. 8 in Materials and Methods, Eq. 15 in Appendix A, and Perkel et al. (1967b)]. The ensemble shift-predictor is flat (Fig. 11C, *thick line*), whereas the shift-predictor for replicate stimuli has a peak (*line with dots*). The peak indicates that the stimulus-locked modulation in neuron 1 and 2 PSTHs (*A, B*, *thin lines*) was correlated. The difference in area between the CCG (*C*, *thin line*) and the two shift-predictors accounts for the difference in r_{CCG} plotted in Figure 10*B* for this pair (*emu035*). In summary, an ensemble shift-predictor fails to capture correlated stimulus-locked modulation, so subtracting it from the raw CCG yields an overestimate of the correlation if correlated stimulus-locked modulation existed in the first place. Thus, when there is little stimulus-locked modulation (as was the case for *rt068* and *rt072* in Fig. 10*B*) or when modulation is present but largely uncorrelated (e.g., *emu034*), using an ensemble shift-predictor is acceptable.

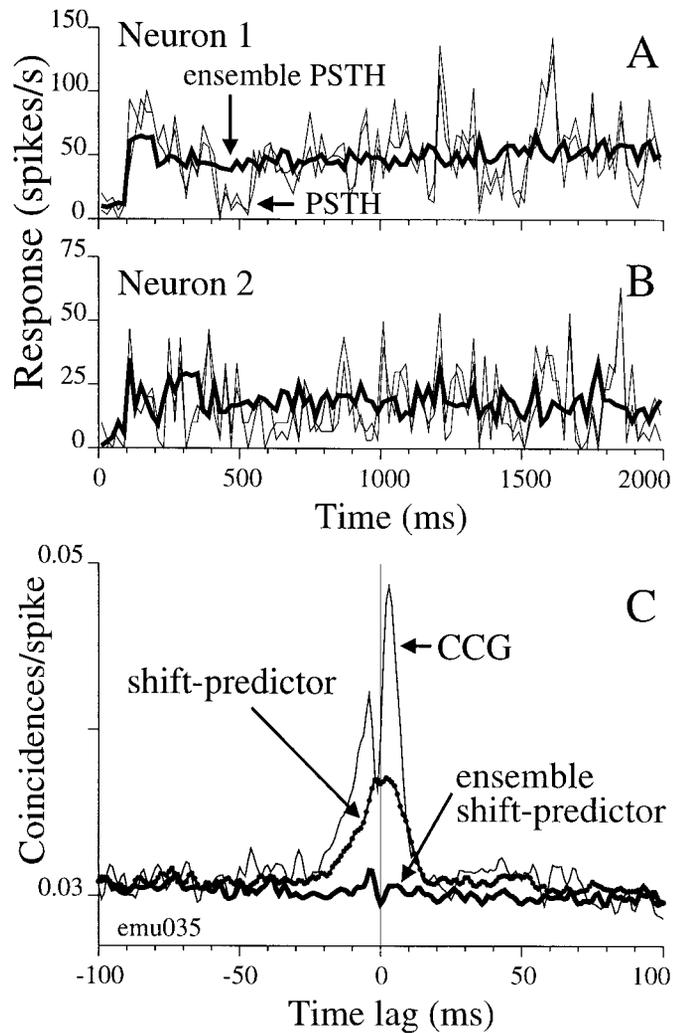


Figure 11. Comparing PSTHs and shift-predictors for responses to ensemble and replicate stimuli for pair *emu035*. *A*, PSTHs (bin size 20 msec) for neuron 1 averaged across 60 trials of different 0% coherence stimuli (ensemble PSTH) and averaged across 30 trials of one particular 0% coherence stimulus (*thin lines*, replicate stimuli, broken into 2 groups of 15 trials) to demonstrate that the modulation is reproducible). *B*, Similar to *A*, but for the simultaneously recorded responses of neuron 2. *C*, The raw CCG (without the shift-predictor subtracted; *thin line*) is plotted for comparison against the ensemble shift-predictor (*thick line*) and the actual shift-predictor (*points*) computed from responses to replicate stimuli. Plots show averages across 13 stimulus conditions ranging from 0 to 51.2% coherence, preferred and null directions. The shift-predictor accounts for roughly half of the area of the CCG peak. It is worth noting that the raw CCG for ensemble stimuli does not differ on average from that for replicate stimuli because both result from cross-correlation of simultaneous responses to stimuli with the same underlying statistics; therefore, only one trace marked *CCG* is shown here. Of course, any particular CCG from repeats of one replicate stimulus will deviate from the average ensemble CCG, but if raw CCGs from many different replicate stimulus sets are averaged together, they will approach the raw ensemble CCG. This is not true for shift-predictors, as seen here, because they are based on responses from non-simultaneous trials.

But for *emu035*, it caused an overestimate of the CCG peak area and of r_{noise} .

One method for estimating the inflation of r_{noise} caused by stochastic stimuli across our database is to compare results for 0 and 100% coherence stimuli. Such a comparison is useful because there is little or no stimulus-locked response modulation for $c =$

100% stimuli, whereas modulation is strong at $c = 0\%$ (Bair and Koch, 1996). For the 20 pairs that we tested at both $c = 0$ and 100% and that consisted of two directional neurons with $\Delta_{\text{PD}} < 90^\circ$, r_{CCG} was on average 0.20 (SD 0.13) at 0% coherence and 0.17 (SD 0.13) at 100%. This 15% decrease is consistent with our hypothesis but was not statistically significant (paired t test; $t = 1.45$; $p = 0.16$). A similar, but unpaired, comparison can be made from the plot of r_{CCG} in Figure 8*D*, which shows a 27% reduction from $c = 0$ to $c = 100\%$ (preferred direction only). Again, this change was not statistically significant (t test; $t = 1.48$; $p = 0.16$). A broader unpaired comparison between all data from discrimination experiments (the vast majority of which was collected at low coherence) and all data from direction tuning experiments (where $c = 100\%$) for pairs with $\Delta_{\text{PD}} < 90^\circ$ showed only an 8% decline in r_{CCG} (32) for the direction tuning data set. These results suggest that inflation of r_{noise} caused by stochastic stimuli is modest across our database.

Finally, we used a simulation to estimate the inflation of r_{noise} caused by the modulated drive resulting from stochastic stimuli. The stimulus drive consisted of randomly occurring bursts of stimulation that simulated those caused by the random occurrence of coherent dots in our motion stimulus. Parameters for the strength and frequency of occurrence of the random bursts determined the amount of trial-to-trial variability and thus the value of r_{noise} . The details of the simulation and a solution for r_{noise} for all parameter values are given in Appendix B. An example of the drive provided by a simulated stimulus during a 1 sec epoch from one trial is shown in Figure 12*A*. The *trace* represents the PSTH for both neurons, which are defined to be identical. For a set of trials governed by the same statistics that generated the trace in *A* (see legend for parameters), the expected value of r_{noise} is 0.04. For a simulation with stronger modulation (*B*), the expected value of r_{noise} is higher, 0.24. Figure 12*D* plots the value of r_{noise} for a wide range of parameter combinations and shows (with *white dots*) the parameters used to generate traces for the examples just described. A comparison of the PSTH for a simulated pair of neurons (*B*) with the measured PSTHs (*C*) for the pair of neurons from Figure 11 reveals a critical difference: the neuronal PSTHs are not identical. This was true although this pair of neurons was as closely matched in preferred direction and bandwidth of direction tuning as any in our database ($\Delta_{\text{PD}} = 9^\circ$; $r_{\text{signal}} = 0.97$). Because nearby neurons have responses that differ in fine detail (DeAngelis et al., 1999), our simulation provides an upper bound on the strength of correlation induced by stochastic stimuli. Furthermore, gauged by responses to replicate stimuli here and in a previous study of MT (Bair and Koch, 1996), the strength of modulation in Figure 12*A* appears typical or above average, whereas that in *B* represents an upper limit to what has been observed. Therefore, our simulations suggest that stochastic stimuli are not likely to inflate r_{noise} by more than ~ 0.04 units on average.

In summary, stochastic stimuli probably inflate our estimates of r_{noise} but cannot be responsible for more than a small fraction of the correlation that we measured. Experimental controls, simulations, and comparisons of incoherent to coherent stimuli suggest that this inflation is likely to range from negligible to at most 20% of our average r_{noise} estimates.

Response variance caused by eye movements

One final potential source of error in our estimate of r_{noise} is the movement of the monkey's eyes. Small saccades executed during fixation could cause correlated signals in neurons with similar

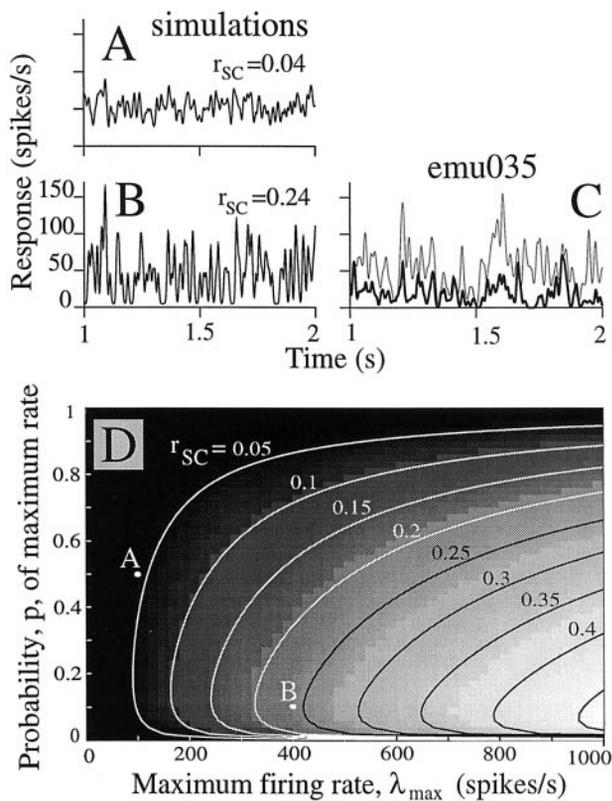


Figure 12. Modeling the interneuronal correlation caused by stimulus variation for a pair of identical neurons. *A*, Simulated time-varying instantaneous mean firing rate for $p = 0.5$, $\lambda_{\max} = 100$ spikes/sec, $\lambda_{\min} = 5$ spikes/sec (smoothed with Gaussian SD 4 msec to achieve a realistic temporal resolution). The *single line* represents the PSTHs for two identical neurons in an ideal pair. See Appendix B for a description of the model. *B*, Similar to *A*, but $p = 0.1$, $\lambda_{\max} = 400$ spikes/sec. *C*, PSTHs for pair *emu035* for comparison to the simulations. The PSTHs for neuron 1 (*thin line*) and neuron 2 (*thick line*) represent a segment of the same data shown by the pairs of *thin lines* in Figure 11, *A* and *B*, but are smoothed like the simulation traces in *A* and *B* here. *D*, For our model, r_{SC} is plotted as a function of all values of p and λ_{\max} (see Eq. 39). Points *A* and *B* mark parameters used to generate traces in *A* and *B*. *Black shading* indicates low correlation, *white* indicates high correlation. Values are given for the *contour lines*.

direction preferences. The potential strength of this effect has been estimated from the influence of eye movements on single-unit MT data (Bair and O'Keefe, 1998), and it was concluded that fixational saccades are too brief and typically too infrequent to create substantial correlation except when occurring on a background of very low firing rate. We found no indication that r_{noise} was higher at lower firing rates (Fig. 8C) and believe that eye movements did not substantially affect estimates of correlation strength in this study.

DISCUSSION

Summary

We have investigated the time scale at which interneuronal correlation arises for pairs of nearby cortical neurons and have explored the relationship between interneuronal correlation and behavioral and stimulus parameters in area MT.

We found that synchrony, revealed by CCG peaks, was closely linked to correlated variability, r_{noise} , at the time scale of the trial. In principle, these two phenomena need not be related (van Kan et al., 1985), but several observations showed that they were

related for our MT pairs. First, the predominant time scale of interneuronal correlation was on the order of 10–100 msec, consistent with numerous cross-correlation studies throughout the visual system of both cat and monkey (Mastrorarde, 1983b; Michalski et al., 1983; Ts'o et al., 1986; Krüger and Aiple, 1988; Nelson et al., 1992; Cardoso de Oliveira et al., 1997) and in auditory cortex (Dickson and Gerstein, 1974; Abeles, 1982; Eggmont and Smith, 1996). Next, CCG peaks at this time scale (10–100 msec) were strikingly predictive of r_{noise} for the behavioral epoch. Although r_{noise} is mathematically related to the total area under the CCG, such a result need not apply to the central CCG region alone. For example, pairs could have had central CCG peaks that were canceled by negative side-lobes, or they could have had excess area distributed across the entire CCG. Neither of these are consistent with our findings. Finally, slow drifts in the gain of single neuronal responses occurred but were not on average correlated between neurons and therefore had little impact on r_{noise} . This result was somewhat surprising because it has been suggested that long-term cross-correlation is common for neurons in primary visual cortex (Bach and Krüger, 1986). Also, because nearby cortical neurons share a large fraction of their inputs, it is unclear how one cell can undergo gain changes that are independent from those of its neighbors. However, if mechanical instability of the electrode in the tissue was the source of the long-term gain changes, it is conceivable that nearby neurons could be affected independently.

In the second part of this study, we found that synchronous activation in pairs of neurons was not related to the monkey's decision on the direction discrimination task and that synchrony was not stronger for perceptually more salient or unified stimuli. Synchrony did not depend on whether the monkey was actively discriminating or passively fixating during stimulus presentation. Finally, the strength of synchrony was similar with and without the stimulus, and it showed little systematic variation with firing rate. We are unable to corroborate reports that synchrony in MT changes with the unity of the stimulus (Kreiter and Singer, 1996; Castelo-Branco et al., 2000) or is nearly abolished during visual stimulation (Cardoso de Oliveira et al., 1997). Experiments using more diverse stimulus configurations will have to resolve these differences. Other studies have suggested that synchrony could signal behavioral events in frontal cortex (Vaadia et al., 1995), encode tone frequency in auditory cortex (deCharms and Merzenich, 1996), indicate attentional selection in somatosensory cortex (Steinmetz et al., 2000), or be involved in arousal, attention, or learning in sensorimotor cortex (Murthy and Fetz, 1996). In contrast, our results portray synchrony and correlation as relatively constant for a typical pair of MT neurons.

In the course of this analysis, we derived two metrics that are useful for determining the strength and time scale of correlation. The TCC provides a systematic way to extract short- and long-term components of the traditional interneuronal correlation coefficient, r_{SC} , for trial-based data, whereas $r_{CCG}(\tau)$ offers an estimate of r_{noise} with lower variance than r_{SC} when the time scale of correlation is shorter than the period during which spikes are counted. We believe that these techniques are potentially useful for comparing correlation across a wide range of data.

Other studies of r_{signal} , r_{noise} , and the CCG

Previous studies of visual cortex have examined r_{noise} , r_{signal} , and spike train CCGs (Gawne and Richmond, 1993; Gawne et al., 1996). They reported r^2 values, interpreted as percentage of explained variance, so we squared our r_{noise} and r_{signal} values

(before averaging) for comparison. Their value of r_{noise}^2 , $\sim 5\%$ for both inferotemporal cortex (IT) and primary visual cortex (V1), was similar to our values: 4.5% for all pairs and 6.4% for directional pairs with $\Delta_{\text{PD}} < 90^\circ$. They found r_{signal}^2 to be 19% in IT and V1 using static, spatial (Walsh) patterns, but this increased to 40% in V1 for conventional bar stimuli. The latter value was comparable to our mean, 48%, for MT. In spite of some similarity between our results, including the fact that over half of their CCGs had significant peaks, they did not comment on the relationship between r_{noise} and the CCG and concluded that the r_{signal} and the CCG were unrelated (they found r_{signal} to be lower for pairs with CCG peaks in IT, but the result failed a significance test). This outcome is different from that depicted in our Figure 2C, which shows a clear relationship between r_{signal} and r_{noise} , where r_{noise} , being r_{CCG} , is a strong reflection of the CCG peak. It seems likely that a relationship like this must exist between r_{signal} and the CCG in both IT and V1 because one consistent feature of CCGs from diverse regions of cortex is that peaks are more common between nearby neurons, particularly within distances associated with cortical columns (Fetz et al., 1991). Cortical columns are clusters of neurons with similar preferences, and such similarity is what r_{signal} , in principle, measures. Maybe differences in the number of cells tested or in the method of estimating the strength of CCG peaks or r_{signal} led to the differences between our results and those of Gawne and collaborators (Gawne and Richmond, 1993; Gawne et al., 1996). For example, the relationship between two-dimensional Walsh patterns and the columnar structure in IT (Fujita et al., 1992; Tanaka, 1996) may be somehow fundamentally different than that between moving patterns and direction columns in MT (Albright et al., 1984).

Consistent with our findings, Bach and Krüger (1986) noted that excess area in the CCG (± 30 msec) was slightly larger for pairs of V1 neurons with strong common variability (i.e., r_{noise}). Also, for both motor and parietal cortex, Lee et al. (1998) found that r_{signal} and r_{noise} were higher for pairs with significant central CCG peaks. All of these results are consistent with the simple notion that sources of common input arrive onto nearby neurons through one or more synapses and thereby create common noise, central peaks in CCGs, and similar tuning curves in pairs of neurons (Shadlen and Newsome, 1998).

Stimulus variance

A major goal of the study from which the present paired MT data arose (Zohary et al., 1994) was to estimate accurately the strength of noise correlation for nearby MT neurons but to do so when those neurons were generating signals that would underlie a psychophysical judgment made by the monkey. The latter constraint led to the use of stochastic stimuli to prevent the monkeys from associating particular stimulus patterns with a reward. In principle, however, stochastic stimuli can bias estimates of r_{noise} upward, as demonstrated by our simulations. We attempted to estimate this bias by comparing responses for replicate and ensemble stimuli, by comparing $c = 0\%$ with $c = 100\%$ data, and by simulating the effect of stochastic stimuli on neuronal responses. The results suggested that the actual r_{noise} value for pairs with similar direction tuning was somewhat less than the measured value of 0.21, but probably not by $>20\%$.

Implications for pooling

Interneuronal correlation places limits on the effectiveness of signal pooling (Johnson et al., 1973; for review, see Parker and Newsome, 1998). Our previous studies showed that the signal-to-

noise ratio (SNR) for a pooled signal was sensitive to even modest values of r_{noise} (Zohary et al., 1994; Shadlen et al., 1996). We can now use our estimates of the time scale of interneuronal correlation to understand how r_{noise} and SNR change with the length of the time window, T , in which signals are pooled.

We simulated pools of spike trains with correlation on the time scale typical for MT (see Fig. 7A legend for methods) and computed the SNR as in Zohary et al. (1994). The SNR for the pooled signal is the expected value, μ_{Σ} , of the sum of spikes from all neurons divided by the SD, σ_{Σ} , of that sum, i.e.:

$$\text{SNR} = \frac{\mu_{\Sigma}}{\sigma_{\Sigma}} = \frac{\mu/\sigma}{\sqrt{\frac{1}{N} + \frac{N-1}{N} r_{\text{noise}}}}, \quad (10)$$

where μ and σ are the mean and SD for spike count from a single neuron, and N is the number of neurons in the pool. Our simulated data were Poisson, so $\sigma^2 = \mu$ and doubling T would increase the SNR by a factor of $\sqrt{2}$ if r_{noise} remained constant, but because correlation was spread over time (Fig. 13A, *thick line*), r_{noise} was lower for shorter T (B, *thick line*). Thus the SNR (Eq. 10) was enhanced for larger pools of neurons at shorter integration times, as shown in C (*thick curves* are squeezed upward in the *bottom right corner*; see legend for details).

Therefore, the time scale of correlation must be taken into account when signals are pooled in short time windows. This may be of relevance to the visual system, where it is likely that some processes underlying visual discrimination operate with integration times from 10 to 100 msec (Oram and Perrett, 1992; Thorpe et al., 1996; Corthout et al., 1999). Here we have focused on one particular pooling model that involves averaging across redundant signals (Zohary et al., 1994; Shadlen et al., 1996; Shadlen and Newsome, 1998). The ultimate role of interneuronal correlation in computations underlying perceptual decisions will depend on details of the actual mechanisms that have yet to be worked out.

APPENDIX A: RELATING SPIKE COUNT CORRELATION TO SPIKE TRAIN CORRELATION

Here we derive an expression that relates the correlation coefficient of spike count, r_{SC} , to the area under the CCG and the ACGs for a set of paired spike trains. A similar relationship was noted earlier by Haim Sompolinsky (personal communication of unpublished notes of 1992 entitled “Statistics of spike counts and spike trains in a stationary process,” pp 1–6), and recently Brody (1999) has noted the relationship between spike count covariance and the area of the CCG, not involving the ACGs. On the basis of our derivation, we propose a metric, $r_{\text{CCG}}(\tau)$, which can provide a lower variance estimate of r_{noise} when interneuronal correlation is limited to a time scale shorter than the trial.

Spike trains from M trials for the two neurons are represented as discrete binary signals of period T at the millisecond resolution, i.e.:

$$x_k^i(t) = \begin{cases} 1, & \text{if on trial } i \text{ neuron } k \text{ fires an action potential} \\ & \text{during the } t^{\text{th}} \text{ millisecond,} \\ 0, & \text{otherwise,} \end{cases} \quad (11)$$

where $k = 1, 2$ and $1 \leq t \leq T$ and $1 \leq i \leq M$. The spike counts for the i^{th} trial are:

$$N_k^i = \sum_{t=1}^T x_k^i(t), \quad (12)$$

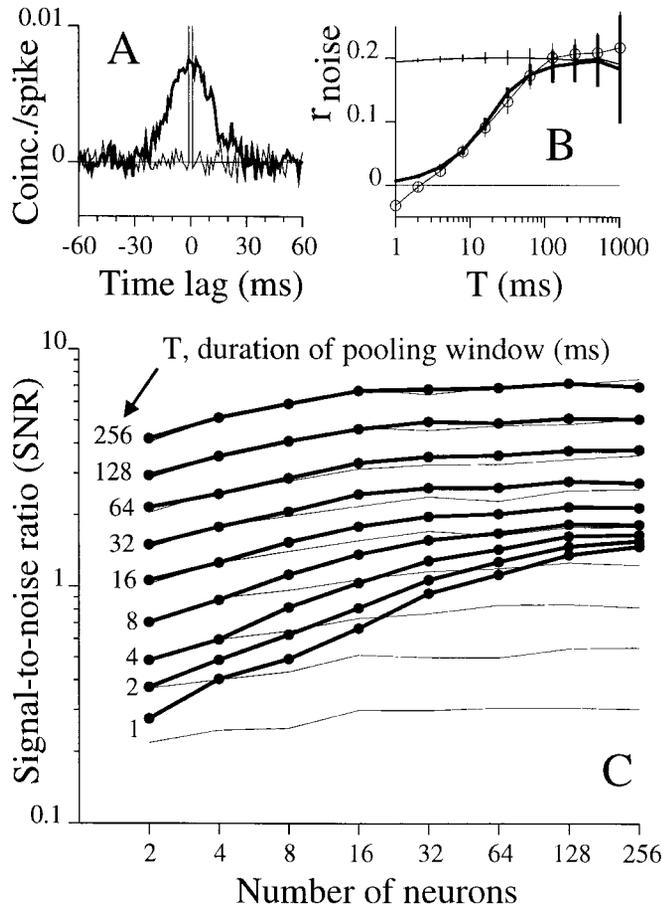


Figure 13. Changes in r_{noise} and SNR as a function of pooling time, T . *A*, CCGs of simulated spike trains are plotted for two hypothetical neuronal populations, one with a realistic time scale of pair-wise correlation (*thick line*, approximates a Gaussian of SD $8\sqrt{2}$ msec) and one with instantaneous correlation (*thin line*, peak at zero is truncated). Simulation method is described in the legend for Figure 7*A*. *B*, For pairs from the two hypothetical populations, r_{noise} was computed as a function of T , the period in which spikes would be counted to form a population response. For instantaneous correlation, r_{noise} was constant (here 0.2) for all integration times (*thin line*). However, for broad correlation (*thick line*), r_{noise} was near zero for short T and increased to the veridical value as T became large relative to the time scale of correlation. Results for the simulated broad correlation were comparable to those for our neuronal data (*open circles*; r_{noise} averaged across 29 neuronal pairs that were directional and had $\Delta_{\text{PD}} < 90^\circ$, coherence series data). The negative value at 1 msec for the neuronal data results from limitations in recording two nearly simultaneous action potentials using one electrode. Error bars for the model show SD across 10 blocks of 200 trials (mean firing rate 40 spikes per second). Error bars are smaller for smaller window sizes because, for example, there are 1000 $T = 1$ msec windows for each $T = 1000$ msec window. *C*, Pooled signals (sums of spike counts) from the hypothetical populations were compared in terms of their SNR (Eq. 10) as a function of neuronal pool size for time windows of various duration. For instantaneous correlation, the SNR curves (*thin lines*) had the same shape for all T but were scaled by $\sqrt{2}$ when T doubled. For the more realistic case of broad correlation, however, the SNR curves (*thick lines*) increased more steeply with pool size for short T because r_{noise} was less for short T (as shown by the *thick line* in *B*).

and the post-stimulus time histograms are:

$$P_k(t) = \frac{1}{M} \sum_{i=1}^M x_k^i(t). \quad (13)$$

The spike train auto-correlation and cross-correlation functions are defined as:

$$C_{jk}(\tau) = \frac{1}{M} \sum_{i=1}^M \sum_{t=1}^T x_j^i(t) x_k^i(t + \tau), \quad (14)$$

where $j = k$ for an auto-correlation and $j = 1, k = 2$ for the cross-correlation, $C_{12}(\tau)$, between neurons 1 and 2. The auto-correlation and cross-correlation of the PSTHs are:

$$S_{jk}(\tau) = \sum_{t=1}^T P_j(t) P_k(t + \tau). \quad (15)$$

For convenience in defining the correlation functions above, we have allowed the time index ($t + \tau$) to take values outside $[1, T]$; therefore, we define $x_k(t)$ and $P_k(t)$ to be zero for $t < 1$ and $t > T$. The function S_{jk} will be referred to as the shift-predictor for the purposes of this appendix because it approximates that portion of the correlation that results from modulation in the PSTHs (Perkel et al., 1967b).

The equation for the correlation coefficient of spike counts:

$$r_{\text{SC}} = \frac{E[N_1 N_2] - E N_1 E N_2}{\sigma_1 \sigma_2}, \quad (16)$$

where E is expected value and σ_k^2 is the variance of the spike count computed over trials, can be rewritten in terms of the cross-correlation equations above. First, observe that:

$$\sigma_k^2 = E N_k^2 - E^2 N_k \quad (17)$$

$$= \frac{1}{M} \sum_{i=1}^M \left[\sum_{t_1=1}^T x_k^i(t_1) \sum_{t_2=1}^T x_k^i(t_2) \right] - \left[\frac{1}{M} \sum_{i=1}^M \sum_{t_1=1}^T x_k^i(t_1) \right] \left[\frac{1}{M} \sum_{i=1}^M \sum_{t_2=1}^T x_k^i(t_2) \right] \quad (18)$$

$$= \frac{1}{M} \sum_{i=1}^M \left[\sum_{\tau=-T}^T \sum_{t=1}^T x_k^i(t) x_k^i(t + \tau) \right] - \sum_{\tau=-T}^T \sum_{t=1}^T P_k(t) P_k(t + \tau) \quad (19)$$

$$= \sum_{\tau=-T}^T [C_{kk}(\tau) - S_{kk}(\tau)]. \quad (20)$$

A similar result holds for the numerator of Equation 16:

$$E[N_1 N_2] - E N_1 E N_2 \quad (21)$$

$$= \frac{1}{M} \sum_{i=1}^M \left[\sum_{t_1=1}^T x_1^i(t_1) \sum_{t_2=1}^T x_2^i(t_2) \right]$$

$$- \left[\frac{1}{M} \sum_{i=1}^M \sum_{t_1=1}^T x_1^i(t_1) \right] \left[\frac{1}{M} \sum_{i=1}^M \sum_{t_2=1}^T x_2^i(t_2) \right] \quad (22)$$

$$= \frac{1}{M} \sum_{i=1}^M \left[\sum_{\tau=-T}^T \sum_{t=1}^T x_1^i(t) x_2^i(t + \tau) \right] - \sum_{\tau=-T}^T \sum_{t=1}^T P_1(t) P_2(t + \tau) \quad (23)$$

$$= \sum_{\tau=-T}^T [C_{12}(\tau) - S_{12}(\tau)]. \quad (24)$$

The following generic expression:

$$A_{jk}(T) = \sum_{\tau=-T}^T [C_{jk}(\tau) - S_{jk}(\tau)], \quad (25)$$

defines the area under the auto- and cross-correlation integrated from $-T$ to T (after the shift-predictor is subtracted). We can rewrite the expression for the correlation coefficient in terms of these areas as follows:

$$r_{SC} = \frac{A_{12}(T)}{\sqrt{A_{11}(T)A_{22}(T)}}. \quad (26)$$

We now define a metric:

$$r_{CCG}(\tau) = \frac{A_{12}(\tau)}{\sqrt{A_{11}(\tau)A_{22}(\tau)}} \quad 1 \leq \tau \leq T, \quad (27)$$

which will be used to estimate the inter-neuronal correlation coefficient by integrating a limited central region of the CCG and ACGs. This measure is equal to the traditional measure, r_{SC} , when $\tau = T$, i.e.:

$$r_{CCG}(T) = r_{SC}. \quad (28)$$

In Results, neuronal data and simulated data are used to demonstrate that $r_{CCG}(\tau)$ can provide a lower variance estimate of r_{noise} .

APPENDIX B: COMPUTING r_{SC} WHEN STIMULUS STRENGTH VARIES

Here we derive an expression for r_{SC} , thus r_{noise} , for a pair of simulated spike trains that arise otherwise independently (i.e., with no common noise) generated from a common stimulus that varies in strength from trial to trial.

Let $f^i(t)$ be the mean firing rate on the i^{th} trial as a function of time (e.g., Fig. 12A), and let two spike trains be generated as independent realizations of an inhomogeneous Poisson processes according to $f^i(t)$. Assume that $f^i(t)$ varies across trial number, i , in such a way that the time-averaged firing rate, λ_i , for any trial has mean μ_λ , variance σ_λ , and probability density g_λ . To derive the correlation in spike count induced by the trial-to-trial changes in $f^i(t)$, we need only consider the statistics of the mean rate, λ , and not the details of the modulation of $f^i(t)$ during the trial. In particular, to compute the correlation coefficient r_{SC} between the spike counts N_1 and N_2 across trials, we must find the expected values and variances required by Equation 16. The expected value of the product of the spike counts can be computed as follows:

$$E[N_1 N_2] = \sum_{j=0}^{\infty} \sum_{k=0}^{\infty} jk \Pr\{N_1 = j, N_2 = k\} \quad (29)$$

$$= \sum_{j=0}^{\infty} \sum_{k=0}^{\infty} jk \int_0^{\infty} g_\lambda(\gamma) \Pr\{N_1 = j, N_2 = k | \lambda = \gamma\} d\gamma \quad (30)$$

$$= \int_0^{\infty} g_\lambda(\gamma) \sum_{j=0}^{\infty} j \Pr\{N_1 = j | \lambda = \gamma\} \sum_{k=0}^{\infty} k \Pr\{N_2 = k | \lambda = \gamma\} d\gamma \quad (31)$$

$$= T^2 \int_0^{\infty} \gamma^2 g_\lambda(\gamma) d\gamma \quad (32)$$

$$= T^2 (\mu_\lambda^2 + \sigma_\lambda^2), \quad (33)$$

where T is the duration of the trial. A derivation similar to that above, but substituting N_1 for N_2 or vice versa, leads to:

$$E[N_1 N_1] = E[N_2 N_2] = T\mu_\lambda + T^2(\mu_\lambda^2 + \sigma_\lambda^2), \quad (34)$$

and a similar but even simpler derivation yields:

$$EN_1 = EN_2 = T\mu_\lambda. \quad (35)$$

Using the identity $\text{VAR}x = E x^2 - E^2 x$ and substituting the results of Equations 33, 34, and 35 into the equation for the correlation coefficient (Eq. 16), we arrive at:

$$r_{SC} = \frac{T\sigma_\lambda^2}{\mu_\lambda + T\sigma_\lambda^2} = \frac{\sigma_N^2}{\mu_N + \sigma_N^2}, \quad (36)$$

where $\mu_N = T\mu_\lambda$ and $\sigma_N^2 = T^2\sigma_\lambda^2$ are used to express the results in terms of spike counts rather than mean rates. This equation states that our simulated spike trains have uncorrelated counts ($r_{SC} = 0$) when there is no trial-to-trial variation in the stimulus strength, i.e., when $\sigma_N^2 = 0$.

To determine the values of μ_N and σ_N^2 , we must define the rate function, $f^i(t)$. Many statistical descriptions are possible, but we chose one that provided modulation which was qualitatively similar to that observed in PSTHs analyzed in our previous study (Bair and Koch, 1996) of responses to replicate stimuli collected under stimulus conditions similar to those of the present study. The rate function, defined as a discrete signal at the resolution of 1 msec, was described by three parameters, a spontaneous firing rate, λ_{\min} , a stimulated firing rate λ_{\max} , and a probability, p , that at each millisecond $f^i(t) = \lambda_{\max}$ (otherwise, $f^i(t) = \lambda_{\min}$). Because for any Bernoulli random variable, X , $E[X] = p$ and $\text{VAR}[X] = pq$ (where p is the probability of success and $q = 1 - p$), it follows that the mean and variance of the trial spike count generated by $f^i(t)$ for trials of duration T seconds are:

$$\mu_N = T[p\lambda_{\max} + (1 - p)\lambda_{\min}], \quad (37)$$

$$\sigma_N^2 = \delta T p (1 - p) (\lambda_{\max} - \lambda_{\min})^2, \quad (38)$$

where λ_{\min} and λ_{\max} are given in spikes per second and $\delta = 0.001$ sec. Substituting this into Equation 36 yields:

$$r_{sc} = \left(\frac{p\lambda_{\max} + (1-p)\lambda_{\min}}{\delta p(1-p)(\lambda_{\max} - \lambda_{\min})^2 + 1} \right)^{-1}. \quad (39)$$

This expression represents the strength of artifactual spike count correlation induced by trial-to-trial stimulus variance for a model of paired spike trains designed to be consistent with MT responses to our dynamic dot stimulus. See Figure 12 and the final section of Results for its application.

REFERENCES

- Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101.
- Abeles M (1982) Local cortical circuits: an electrophysiological study. Berlin: Springer-Verlag.
- Aertsen AMHJ, Gerstein GL, Habib MK, Palm G (1989) Dynamics of neuronal firing correlation: modulation of “effective connectivity.” *J Neurophysiol* 61:900–917.
- Albright TD, Desimone R, Gross CG (1984) Columnar organization of directionally selective cells in visual area MT of the macaque. *J Neurophysiol* 51:16–31.
- Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273:1868–1871.
- Bach M, Krüger J (1986) Correlated neuronal variability in monkey visual cortex revealed by a multi-microelectrode. *Exp Brain Res* 61:451–456.
- Bair W, Koch C (1996) Temporal precision of spike trains in extrastriate cortex. *Neural Comput* 8:1185–1202.
- Bair W, O’Keefe LP (1998) The influence of fixational eye movements on the response of neurons in area MT of the macaque. *Vis Neurosci* 15:779–786.
- Bair W, Koch C, Newsome WT, Britten KH (1994) Power spectrum analysis of bursting cells in area MT in the behaving monkey. *J Neurosci* 14:2870–2892.
- Bair W, Zohary E, Koch C (1996) Correlated neuronal response: time scales and mechanisms. In: *Advances in neural information processing systems 8* (Touretzky DS, Mozer MC, Hasselmo ME, eds), pp 68–74. Cambridge, MA: MIT.
- Bair W, Zohary E, Newsome W (1999) Synchrony in cross-correlograms linked to spike count correlation in pairs of neurons in area MT. *Soc Neurosci Abstr* 25:276.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Response of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 10:1157–1169.
- Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13:87–100.
- Brody CD (1998) Slow covariations in neuronal resting potential can lead to artefactually fast cross-correlation in their spike trains. *J Neurophysiol* 80:3345–3351.
- Brody CD (1999) Correlations without synchrony. *Neural Comput* 11:1527–1535.
- Cardoso de Oliveira S, Thiele A, Hoffmann K-P (1997) Synchronization of neuronal activity during stimulus expectation in a direction discrimination task. *J Neurosci* 17:9248–9260.
- Castelo-Branco M, Goebel R, Neuenschwander S, Singer W (2000) Neural synchrony correlates with surface segregation rules. *Nature* 405:685–689.
- Celebrini S, Newsome WT (1995) Microstimulation of extrastriate area MST influences performance on a direction discrimination task. *J Neurophysiol* 73:437–448.
- Corthout E, Uttl B, Walsh V, Hallett M, Cowey A (1999) Timing of activity in early visual cortex as revealed by transcranial magnetic stimulation. *NeuroReport* 10:2631–2634.
- Crist CF, Yamasaki DSG, Komatsu H, Wurtz RH (1988) A grid system and a microsyringe for single cell recording. *J Neurosci Methods* 26:117–122.
- Das A, Gilbert CD (1995) Receptive field expansion in adult visual cortex is linked to dynamic changes in strength of cortical connections. *J Neurophysiol* 74:779–792.
- DeAngelis GC, Newsome WT (1999) Organization of disparity-selective neurons in macaque area MT. *J Neurosci* 19:1398–1415.
- DeAngelis GC, Ghose GM, Ohzawa I, Freeman RD (1999) Functional micro-organization of primary visual cortex: receptive field analysis of nearby neurons. *J Neurosci* 19:4046–4064.
- deCharms RC, Merzenich MM (1996) Primary cortical representation of sounds by the coordination of action-potential timing. *Nature* 381:610–613.
- Dickson JW, Gerstein GL (1974) Interactions between neurons in auditory cortex of the cat. *J Neurophysiol* 37:1239–1261.
- Eggermont JJ, Smith GM (1995) Rate covariance dominates spontaneous cortical unit-pair correlograms. *NeuroReport* 6:2125–2128.
- Eggermont JJ, Smith GM (1996) Neural connectivity only accounts for a small part of neural correlation in auditory cortex. *Exp Brain Res* 110:379–391.
- Evarts EV (1968) Methods for recording activity of individual neurons in moving animals. *Methods Med Res* 11:241–250.
- Fetz E, Toyama K, Smith W (1991) Synaptic interactions between cortical neurons. In: *Normal and altered states of function. Cerebral cortex, Vol 9* (Peters A, Jones EG, eds), pp 1–47. New York: Plenum.
- Fujita I, Tanaka K, Ito M, Chen K (1992) Columns for visual features of objects in monkey inferotemporal cortex. *Nature* 360:343–346.
- Funahashi S, Inoue M (2000) Neuronal interactions related to working memory processes in the primate prefrontal cortex revealed by cross-correlation analysis. *Cereb Cortex* 10:535–551.
- Gawne TJ, Richmond BJ (1993) How independent are the messages carried by adjacent inferior temporal cortical neurons? *J Neurosci* 13:2758–2771.
- Gawne TJ, Kjaer TW, Hertz JA, Richmond BJ (1996) Adjacent visual cortical neurons share about 20% of their stimulus-related information. *Cereb Cortex* 6:482–489.
- Geisler WS, Albrecht DG (1997) Visual cortex neurons in monkeys and cats: detection, discrimination, and identification. *Vis Neurosci* 14:897–919.
- Green DM, Swets JA (1966) Signal detection theory and psychophysics. New York: Wiley.
- Hawken MJ, Shapley RM, Gross DH (1996) Temporal frequency selectivity in monkey visual cortex. *Vis Neurosci* 13:477–492.
- Hoel PG, Port SC, Stone CJ (1971) Introduction to statistical theory. Boston: Houghton Mifflin.
- Johnson KO (1980) Sensory discrimination: neural processes preceding discrimination decision. *J Neurophysiol* 43:1793–1815.
- Johnson KO, Darian-Smith I, LaMotte C (1973) Peripheral neural determinants of temperature discrimination in man: a correlative study of responses to cooling skin. *J Neurophysiol* 36:347–370.
- Judge SJ, Richmond BJ, Chu FC (1980) Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res* 20:535–538.
- Kreiter AK, Singer W (1996) Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *J Neurosci* 16:2381–2396.
- Krüger J, Aiple F (1988) Multimicroelectrode investigation of monkey striate cortex: spike train correlations in the infragranular layers. *J Neurophysiol* 60:798–828.
- Lee D, Port NL, Kruse W, Georgopoulos AP (1998) Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J Neurosci* 18:1161–1170.
- Mastrorade DN (1983a) Correlated firing of cat retinal ganglion cells. I. Spontaneously active inputs to X- and Y-cells. *J Neurophysiol* 49:303–324.
- Mastrorade DN (1983b) Correlated firing of cat retinal ganglion cells. II. Responses of X- and Y-cells to single quantal events. *J Neurophysiol* 49:325–349.
- Maunsell JHR, Van Essen DC (1983) Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *J Neurophysiol* 49:1127–1147.
- Melssen WJ, Epping WJM (1987) Detection and estimation of neural connectivity based on crosscorrelation analysis. *Biol Cybern* 57:403–414.
- Michalski A, Gerstein GL, Czarkowska J, Tarnecki R (1983) Interactions between cat striate cortex neurons. *Exp Brain Res* 51:97–107.
- Moore GP, Segundo JP, Perkel DH, Levitan H (1970) Statistical signs of synaptic interaction in neurons. *Biophys J* 10:876–900.
- Morgan MJ, Ward R (1980) Conditions for motion flow in dynamic visual noise. *Vision Res* 20:431–435.
- Mountcastle VB, Talbot WH, Sakata H, Hyvärinen J (1969) Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. Neuronal periodicity and frequency discrimination. *J Neurophysiol* 32:452–484.
- Murasugi CM, Salzman CD, Newsome WT (1993) Microstimulation in visual area MT: effects of varying pulse amplitude and frequency. *J Neurosci* 13:1719–1729.
- Murthy VN, Fetz EE (1996) Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *J Neurophysiol* 76:3968–3982.
- Nelson JJ, Salin PA, Munk MHJ, Arzi M, Bullier J (1992) Spatial and temporal coherence in cortico-cortical connections: a cross-correlation study in areas 17 and 18 in the cat. *Vis Neurosci* 9:21–37.
- Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. *Nature* 341:52–54.
- Oram MW, Perrett DI (1992) Time course of neural responses discriminating different views of the face and head. *J Neurophysiol* 68:70–84.

- Panzeri S, Schultz SR, Treves A, Rolls ET (1999) Correlations and the encoding of information in the nervous system. *Proc R Soc Lond B Biol Sci* 266:1001–1012.
- Parker AJ, Newsome WT (1998) Sense and the single neuron: probing the physiology of perception. *Annu Rev Neurosci* 21:227–277.
- Perkel DH, Gerstein GL, Moore GP (1967a) Neuronal spike trains and stochastic point processes I. The single spike train. *Biophys J* 7:391–418.
- Perkel DH, Gerstein GL, Moore GP (1967b) Neuronal spike trains and stochastic point processes II. Simultaneous spike trains. *Biophys J* 7:419–440.
- Quick RF (1974) A vector magnitude model of contrast detection. *Kybernetik* 16:65–67.
- Raiguel SE, Xiao D-K, Marcar VL, Orban GA (1999) Response latency of macaque area MT/V5 neurons and its relationship to stimulus parameters. *J Neurophysiol* 82:1944–1956.
- Roelfsema PR (1998) Solutions for the binding problem. *Z Naturforsch [C]* 53:691–715.
- Salzman CD, Newsome WT (1994) Neural mechanisms for forming a perceptual decision. *Science* 264:231–237.
- Salzman CD, Murasugi CM, Britten KH, Newsome WT (1992) Microstimulation in visual area MT: effects on direction discrimination performance. *J Neurosci* 12:2331–2355.
- Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896.
- Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci* 16:1486–1510.
- Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. *Annu Rev Neurosci* 18:555–586.
- Steinmetz PN, Roy A, Fitzgerald PJ, Hsiao PJ, Johnson KO, Niebur E (2000) Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature* 404:187–190.
- Tanaka K (1996) Inferotemporal cortex and object vision. *Annu Rev Neurosci* 19:109–139.
- Thorpe S, Fize D, Marlot C (1996) Speed of processing in the human visual system. *Nature* 381:520–522.
- Ts'o DY, Gilbert CD (1988) The organization of chromatic and spatial interactions in the primate striate cortex. *J Neurosci* 8:1712–1727.
- Ts'o DY, Gilbert CD, Wiesel TN (1986) Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J Neurosci* 6:1160–1170.
- Vaadia E, Haalman I, Abeles M, Bergman H, Prut Y, Slovin H, Aertsen A (1995) Dynamics of neuronal interactions in monkey cortex in relation to behavioral events. *Nature* 373:515–518.
- van Kan PLE, Scobey RP, Gabor AJ (1985) Response covariance in cat visual cortex. *Exp Brain Res* 60:559–563.
- Watson AB (1979) Probability summation over time. *Vision Res* 19:515–522.
- Wörgötter F, Daunicht WJ, Eckmiller R (1986) An on-line spike form discriminator for extracellular recordings based on an analog correlation technique. *J Neurosci Methods* 17:141–151.
- Zeki SM (1974) Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *J Physiol (Lond)* 236:549–573.
- Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370:140–143.