

Measuring and interpreting neuronal correlations

Marlene R Cohen¹ & Adam Kohn²

Mounting evidence suggests that understanding how the brain encodes information and performs computations will require studying the correlations between neurons. The recent advent of recording techniques such as multielectrode arrays and two-photon imaging has made it easier to measure correlations, opening the door for detailed exploration of their properties and contributions to cortical processing. However, studies have reported discrepant findings, providing a confusing picture. Here we briefly review these studies and conduct simulations to explore the influence of several experimental and physiological factors on correlation measurements. Differences in response strength, the time window over which spikes are counted, spike sorting conventions and internal states can all markedly affect measured correlations and systematically bias estimates. Given these complicating factors, we offer guidelines for interpreting correlation data and a discussion of how best to evaluate the effect of correlations on cortical processing.

Understanding how populations of neurons encode information and guide behavior is a major focus of systems neuroscience. Cortical neurons respond with variable strength to repeated presentations of identical stimuli^{1,2}. This variability is often shared among neurons, and such correlations in trial-to-trial responsiveness can have a substantial effect on the amount of information encoded by a neuronal population²⁻⁶. Early studies, for instance, showed that correlations reduce the signal-to-noise of a pooled population response, as shared fluctuations in response cannot be averaged away⁶. More recently, it has been shown that attentional modulation of correlations accounts for more of its effect on sensory coding than modulation of firing rate^{7,8}. Determining how correlations are affected by stimulus drive⁹⁻¹², learning or experience^{10,13,14}, or changes in behavioral context^{7,8,15-17} is therefore likely to be as important as understanding how these factors affect the firing rates of individual cells¹⁸.

Correlations can also provide important information about the functional architecture of neuronal networks. Correlation analysis has been used to infer connectivity in the retina¹⁹, between the visual thalamus and cortex²⁰, and between neurons in cortex^{9,21}. For a given circuit, changes in correlations under different stimulus or behavioral conditions can provide a signature of network function or computations that may be difficult to discern from measurements of individual neuronal firing rates^{10,12,14-18}. For example, when an animal actively explores its environment, responses in sensory cortex are desynchronized, even when sensory input is disrupted and neuronal firing rates are unchanged¹⁶.

Studying correlations is, however, inconvenient. It requires large amounts of data from simultaneously recorded neurons, and raw correlation values can be difficult to interpret. The recent advent of recording techniques such as multielectrode arrays and two-photon

imaging has made obtaining such recordings easier but has also highlighted the need to understand the experimental and physiological factors that can lead to different estimates and conclusions.

Forms of correlation and their basic properties

Correlation is a normalized measure of covariation. It has commonly been used to refer to two distinct phenomena (**Fig. 1**). One use refers to tuning similarity, measured as the correlation in the mean responses of two neurons to an ensemble of stimuli (termed signal correlation or r_{signal} ; see **Fig. 1a,c** and **Box 1**). The second use of correlation, and our focus here, is as a measure of the degree to which trial-to-trial fluctuations in response strength are shared by a pair of neurons. This is typically quantified as the Pearson correlation of the spike count responses to repeated presentations of identical stimuli under the same behavioral conditions (spike count correlation or r_{SC} , also called noise correlation; **Fig. 1a,b** and **Box 1**).

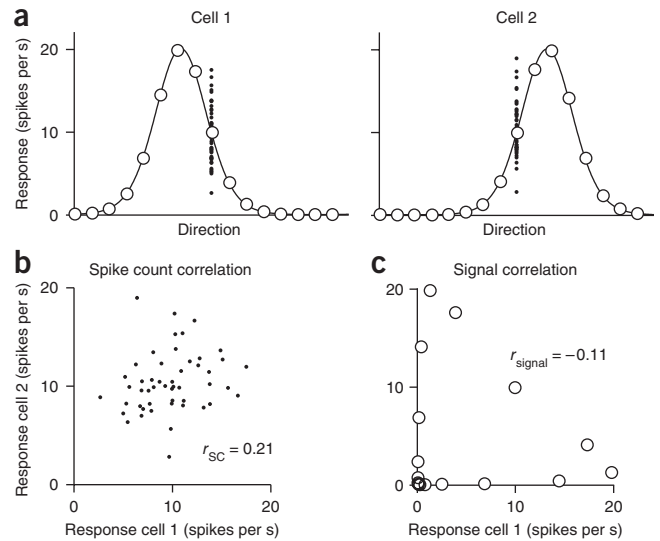
Co-fluctuations in the responses of a pair of neurons can arise over a range of timescales^{8,12,22,23}, from the precise temporal alignment of spikes (that is, synchrony) to slower changes in excitability (**Box 1**). The timescale over which correlated activity affects the responses of downstream neurons is unknown, but membrane time constants suggest that it occurs over tens of milliseconds or less. However, most work on the relationship of spike count correlations to population coding and behavior has been based on responses measured over the duration of a stimulus presentation or behavioral trial (typically hundreds of milliseconds).

Over the past two decades, spike count correlations have been measured in many cortical areas under a variety of behavioral and stimulus conditions (**Table 1**). These studies have reported a range of values, but correlations are typically small and positive. They tend to be highest for pairs of neurons that are near each other²⁴⁻²⁶ and have similar functional properties or tuning (high r_{signal})^{6,7,12,13,15,25-31}. Pairs recorded from opposite hemispheres have very low correlations⁷. These properties suggest that correlations reflect co-fluctuations in the responses of restricted subsets of neurons, rather than global fluctuations that affect all cells.

¹Department of Neurobiology, Harvard Medical School, Boston, Massachusetts, USA. ²Department of Neuroscience and Department of Ophthalmology and Visual Sciences, Albert Einstein College of Medicine, Bronx, New York, USA. Correspondence should be addressed to M.R.C. (cohenm@pitt.edu) or A.K. (adam.kohn@einstein.yu.edu).

Published online 27 June 2011; doi:10.1038/nn.2842

Figure 1 Types of pair-wise neuronal correlations. **(a)** Tuning curves for two hypothetical direction-selective neurons. Open circles show mean responses to different directions of motion and small points show responses to individual presentations of a stimulus at a particular direction. **(b)** Spike count or ‘noise’ correlation (r_{SC}) measures the correlation between fluctuations in responses to the same stimulus. Here, each point represents the response of the two neurons on one presentation of an individual stimulus. **(c)** Signal correlation (r_{signal}) measures the correlation between the two cells’ mean responses to different stimuli. Each point represents the mean response to a given direction of motion. Because the responses of cell 2 increase of a range of motion directions in which the responses of cell 1 decline, signal correlation is negative.



Correlation strength is also likely to depend on local circuitry or architecture. For instance, r_{SC} is weak in the input layers of primary sensory cortex³² (M.A. Smith and A.K., unpublished data; J. Hansen and V. Dragoi (University of Texas—Houston), personal communication). Correlations in motor areas seem to be consistently lower than those in sensory cortex (see **Table 1**).

The influence of distance, tuning similarity and architecture can explain some of the variability across studies, but even studies that sample similarly can arrive at quite different estimates of correlation strength. We show here that many of these discrepancies can be explained by differences in other factors that can systematically bias correlation estimates: namely, response strength, the time period for counting spikes, spike sorting and fluctuations in internal states.

Why is it important to understand the influence of these factors and, more generally, differences in estimates across studies? It is not because the mean strength of correlations is a particularly critical quantity. Even very weak correlations can substantially affect the information encoded by a population of neurons, and the structure of correlations can have a much stronger influence than their mean strength^{2–6}. Furthermore, other properties of the distribution of correlation coefficients, such as its variance, may be more informative about the underlying circuit than the mean³³. However, understanding differences in correlations across brain regions or in different stimulus or task conditions is critical for both elucidating their role in sensory processing and making inferences about the circuitry and mechanisms that generate them. For this reason, we explore the way that different experimental and physiological factors affect measurements of correlations and discuss guidelines for interpreting correlation data.

Experimental factors that affect r_{SC}

Correlations are small when based on few spikes. Correlations in pairs of neurons that fire few spikes per trial are weaker than in pairs that respond more strongly^{7,8,28,34}. Mathematically, correlations between discrete variables such as spike counts need not depend on their magnitude. Binary variables can be perfectly correlated, uncorrelated or correlated to any intermediate degree. The dependence of

r_{SC} on spike count is therefore not a mathematical given but a biological and experimental phenomenon, as described below.

The number of spikes produced by a neuron is determined by its underlying firing rate and the time window over which responses are measured, and both factors vary considerably across studies. Firing rates depend on the stimulus, the animal’s cognitive state and the neurons’ tuning. In some studies, the stimulus is tailored to the tuning preferences of the neurons being studied^{6,12,29}. In others, a common stimulus is used to drive a large number of simultaneously recorded cells^{7,26,28}, yielding weaker responses on average. In addition, neurons in some cortical areas, such as V1 and MT, are easier to drive than those in areas in which stimulus preferences are less understood. The time window used to count spikes is purely an experimental decision, and therefore also varies across studies. Below, we show that both low firing rates and brief measurement windows can lead to lower measured values of r_{SC} .

The spike threshold can reduce spike count correlations. The relationship between r_{SC} and firing rate depends largely on the proportion of subthreshold events that are masked by the spiking threshold³⁴. Correlations in spiking responses arise because of co-fluctuations in synaptic input, which give rise to correlated membrane potential fluctuations in pairs of neurons^{16,33,35–37}. The degree to which shared membrane potential fluctuations are observable in spiking responses depends on the firing rates of the cells; when the mean membrane potential is far below threshold, responses are weak and many of the shared membrane potential fluctuations are unobservable in the spiking responses^{34,37,38}.

Box 1 Types of correlations

Signal correlation (r_{signal}) measures the correlation coefficient between mean responses to different stimuli. This measure is often used to quantify the extent to which a pair of neurons has similar tuning or other functional properties. Decreases in this type of correlation, such as through adaptation⁹¹ or contextual modulation⁹², can lead to sparsening of population responses.

Spike count correlation (r_{SC} , also called noise correlation) is the Pearson’s correlation coefficient of spike count responses to repeated presentations of identical stimuli, under the same behavioral conditions. Spike counts are typically measured over the timescale of a stimulus presentation or a behavioral trial, which range from a few hundred milliseconds to several seconds (see **Table 1**). Spike count correlations are proportional to the integral under the spike train cross-correlogram²².

Synchrony measures the extent to which the timing of spikes is precisely aligned, typically on the timescale of one or a few milliseconds. It is typically quantified using the sharp peak of in the cross-correlogram. Dependencies in the timing of individual spikes can also be measured in the frequency domain, using spike-spike coherence.

Long-timescale correlation (r_{LT}) measures the extent to which a neuron’s response on one trial is correlated with a second neuron’s response on trials in the future or past²². It is used to quantify the influence of slow fluctuations in responsiveness on r_{SC} . When measured, r_{LT} has been found to be close to 0.



Table 1 Summary of studies measuring spike count correlations in primates

Reference number	Area	Firing rate (spikes per s)	Duration (ms)	State (task, anesthesia, etc.)	r_{SC}
12*	V1	~25	2,560	Anesthetized	0.2
26*	V1	~8	1,280	Anesthetized	0.16
23	V1		1,894	Anesthetized	0.25
31	V1			Anesthetized	0.26
13*	V1	~50	1,860	Fixation	0.25
28*	V1	~3	500	Fixation	0.01
82	V1		400	Tracing	0.18
83*	V1	30	1,000	Discrimination	0.1
A. Zandvakili and A.K., unpublished data*	V2	5	1,000	Anesthetized	0.11
M. Smith and M. Sommer (University of Pittsburgh), personal communication*	V4	5.2	1,000	Fixation	0.05
7*	V4	21	200	Attention/detection task	0.04
8*	V4	>5, ~20	800	Attention/tracking task	0.05
A.B.G. Graf (New York University), personal communication*	MT	~10	300	Anesthetized	0.09
29*	MT	~20	500	Fixation	0.1
15*	MT	28.5	500	Discrimination	0.13
6/22*	MT	~20	1,000	Discrimination	0.15
84	Perirhinal	~12	200–500	Fixation/matching task	0.02
85	Supp motor area		66 or 200	Serial reaching	0.013
27	Supp motor area	~15	200	Reaching	0.02
86	Premotor areas	~5	400	Grasping/imagery task	0.02
87	M1	~20	600	Reaching	0.1–0.2
25	Motor/parietal; areas 2/5	~5	1,000	Reaching	0.02–0.04
88	Substantia nigra	58	500	Cue matching	0.01–0.04
89	FEF	~50	A few hundred	Visual search	0.05–0.2
90	FEF	~20	~200	Visual search	0.09
24	Prefrontal	~5	3,000	Delayed saccade task	0.08

These studies measured correlations in a variety of brain areas, behavioral and stimulus conditions, and measurement durations and between pairs of neurons that varied in the cortical distance and tuning similarity. When multiple values of correlations, firing rates and measurement windows were reported, we list either the average or most common value that was listed in the text or estimated from summary figures. Supp, supplementary; FEF, frontal eye field.

This relationship is illustrated with a simple simulation shown in **Figure 2** (similar to ref. 34). We simulated correlated membrane potentials by picking values for each of two neurons from a bivariate Gaussian distribution. The membrane potential produces a spike response defined by the nonlinearity. The shape of the nonlinearity is not critical; ours was chosen so that the variance of the spiking response is nearly the same as the mean for both weak and strong responses³⁹.

We set the membrane potential correlation to 0.2 in all of our simulations, but the measured spike count correlation depended on response strength. When the mean membrane potential is above threshold (**Fig. 2a**), the correlation in spiking responses is similar to that of the subthreshold response. However, when the mean membrane potential is far below threshold (**Fig. 2b**), r_{SC} is markedly lower than between the subthreshold responses. This masking of correlated activity cannot be overcome by making more observations, which reduces the variance of correlation measurements, but does not alter the mean. If membrane potential correlations are the same for stimuli that drive weak and strong responses, the spiking responses for the latter will therefore be more correlated (**Fig. 2c**), reaching asymptote at the strength of the underlying membrane potential correlation.

The dependence of correlation on response strength is typically assessed by comparing r_{SC} to the geometric mean response of the two neurons^{7,8,28,34}. However, r_{SC} will be reduced if either

neuron responds weakly. For example, a pair in which one neuron has a mean response of 0.01 spikes per s and the other 100 spikes per s has the same geometric mean response as a pair whose mean responses are both 1 spike per s, but the measured r_{SC} of the first pair is only 15% of the underlying correlation compared with 85% in the second pair. **Figure 2d** shows r_{SC} as a function of the response strength of the two neurons. A dependence of r_{SC} on the geometric mean response would appear as diagonal stripes from the top left to bottom right. Instead, vertical and horizontal bands are seen, indicating that the magnitude of r_{SC} depends more on the minimum response of the two neurons than their mean.

Counting spikes over short windows can lead to weaker correlation. The number of spikes a neuron fires also depends on the time window over which responses are measured. The studies in **Table 1** use windows that range from tens of milliseconds to multiple seconds. Counting spikes over short epochs can lead to weaker observed values of r_{SC} , even if both neurons are sufficiently responsive to avoid the effect of thresholding described in **Figure 2**.

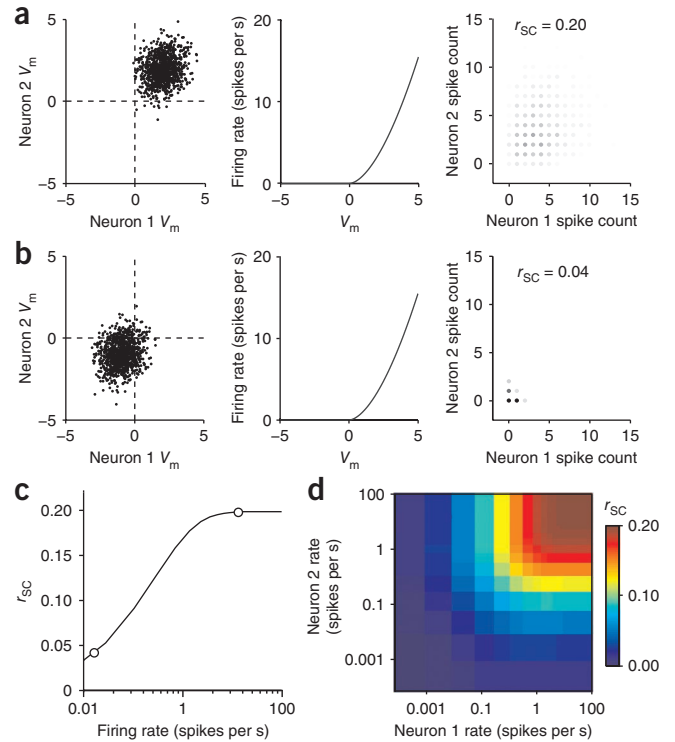
We ran additional simulations to illustrate the dependence of r_{SC} on measurement window (**Fig. 3**). To do so, it was necessary to use a framework in which we could control the timescale of correlation. As the simulations shown in **Figure 2** do not specify a timescale, we instead imposed correlations by adding a small number of common spikes to the otherwise independent (Poisson) spike

trains of two simulated neurons (**Fig. 3a,b**; also see ref. 40).

This simulation illustrates that correlations are systematically underestimated if the counting window is shorter than the jitter in the timing of the common spikes. If the common spikes occur at the same instant (**Fig. 3a**), the resulting synchrony will be evident in spike count correlations based on responses measured over arbitrarily small windows (**Fig. 3c**). If the times of these common spikes are jittered (**Fig. 3b**), however, the resulting spike count correlation will only be fully evident when it is calculated from responses during longer response epochs (**Fig. 3c**). For instance, when spike times are jittered using a Gaussian distribution with an s.d. of 80 ms, a window of several hundred milliseconds is needed to capture the full strength of correlation. A similar dependence on measurement window is observed for Poisson distributed spikes conditioned on correlated underlying firing rates (see ref. 22 and **Supplementary Results** for analytical description). Note that in this scenario correlations do not arise from nearly synchronous spikes, but the observed correlations are still smaller for brief measurement windows.

Several studies have measured the timescale of correlation in cortex, providing estimates ranging from tens²² to a few hundred milliseconds^{8,12,23}. Measurements using response windows briefer than these timescales are thus almost certain to yield weaker correlations than those using longer response windows.

Figure 2 Measured correlations are small when responses are weak. **(a)** We drew intracellular voltage events from a bivariate Gaussian distribution of voltage relative to threshold (V_m) such that the two neurons had the same mean voltage and the correlation coefficient between membrane potentials was 0.2 (left). These voltage events were converted to extracellular firing rates by passing them through a nonlinearity such that the firing rate on the i th trial $f_i = V_{m_i}^{1.7}$ (center). We picked this exponent so that the variance of the output rates was approximately equal to the mean³⁹. We defined the spike count on each trial to be equal to the firing rate rounded to the nearest whole spike (right). In the case of high mean voltages, the measured spike count correlation is close to the input correlation ($r_{SC} = 0.20$). Shading indicates the number of observations. **(b)** Data are presented as in **a** for low mean V_m . When the threshold masks subthreshold events, the measured spike count correlation ($r_{SC} = 0.04$) is much lower than the membrane potential correlation. **(c)** Measured r_{SC} as a function of firing rate, using the simulations in **a** and **b** when the two cells had the same mean rate. The circles represent the correlations and mean rates in **a** and **b**. **(d)** Measured r_{SC} as a function of the firing rates of each of the two cells. We simulated responses using identical methods as those in **a** and **b**, but we allowed the rates of the two neurons in a pair to differ. Correlations depend more strongly on the minimum rate in the pair than the mean.



Effect of spike sorting errors on measured r_{SC} . Issues of recording quality or spike sorting can artifactually increase or decrease measurements of r_{SC} . In general, errors that add independent variability to the responses of one neuron will bias estimates of correlations toward zero, whereas errors that involve combining the responses of multiple cells will increase the magnitude of measured r_{SC} .

One spike sorting error that can lead to consistent overestimation of correlations is mistaking multiunit activity as spikes from a single neuron. Combining several units into one effectively averages out variability that is independent of each cell, so the correlation between two clusters of multiunit activity will be larger than between pairings of the constituent neurons. To quantify this influence, we simulated multiunit activity by grouping the responses of individual neurons produced from the simulation in **Figure 2**. We computed r_{SC} between

increasingly large groupings of neurons, made by simply summing the responses of individual units whose pair-wise correlations ranged from very small (0.001) to more typical (0.1) values (**Fig. 4**).

These simulations show that, when pair-wise correlations are weak, the measured value of r_{SC} grows slowly with the number of units grouped together. Specifically, the value of r_{SC} between multiunit clusters ($r_{SC\text{-measured}}$) is given by

$$r_{SC\text{-measured}} = \frac{nr_{SC\text{-pair}}}{(n-1)r_{SC\text{-pair}} + 1} \quad (1)$$

where n is the number of units grouped together and $r_{SC\text{-pair}}$ is the pair-wise correlation, which in this simulation was the same between all units, whether they were in the same or different groupings^{41,42}. For $nr_{SC\text{-pair}} \ll 1$, r_{SC} thus grows linearly with the number of units n contributing to the multiunit clusters (that is, as $nr_{SC\text{-pair}}$). If, for example, the underlying pair-wise correlation were 0.01, one would need to record simultaneously from clusters of nearly 20 cells to obtain the r_{SC} of 0.2 that has been reported in many studies. Recordings from such large groups of cells would be evident by a proportional increase in firing rate and thus easy to distinguish from single-unit activity.

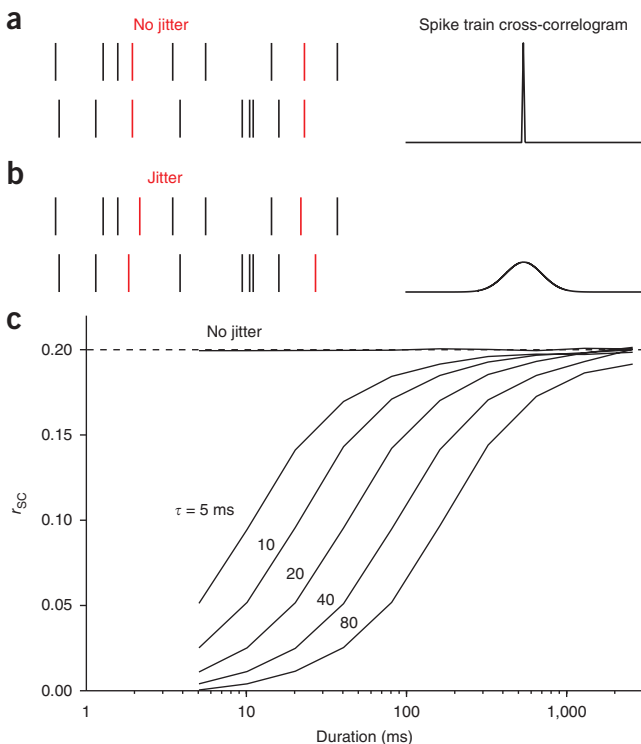
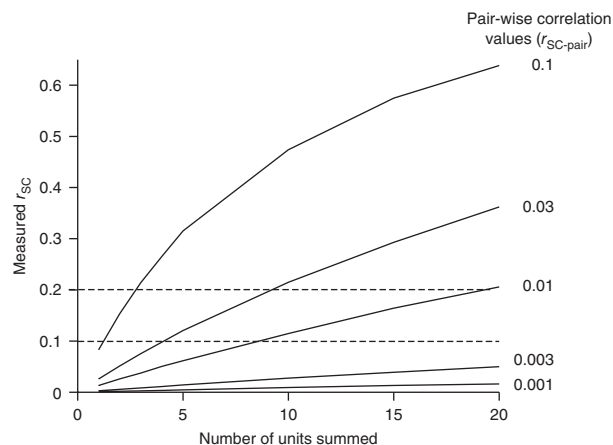


Figure 3 Counting spikes over short response windows can decrease measured correlations. **(a)** We simulated correlated spike trains as a combination of independent Poisson spike trains (black spikes in the schematic on the left, mean = 20 spikes per s) and inserted shared, synchronous spikes (red spikes, mean = 5 spikes per s). Thus, each neuron's response was the sum of independent and shared spikes (mean = 25 spikes per s). When there is no jitter in the timing of the shared spikes, the cross-correlogram has a sharp peak at 0-ms time lag (right). **(b)** Data are presented as in **a** for jittered spike times. We simulated variable timescales by jittering the timing of the correlated spikes by an amount picked from a Gaussian distribution (left, red spikes). This results in a cross-correlogram with a peak whose width depends on the s.d. of the Gaussian distribution (right). **(c)** Measured r_{SC} as a function of counting window for several timescales of correlation. The number next to each curve corresponds to the s.d. of the Gaussian jitter in milliseconds.

Figure 4 Measured correlations grow slowly with the number of units that contribute to multiunit activity. We simulated single unit activity by choosing random variables from a bivariate Gaussian with covariance defined to give pair-wise correlation values of 0.001–0.1. We then computed r_{SC} between sums of these variables (as one would do when recording multiunit activity). Measured r_{SC} increases with the number of ‘units’ contributing to the multiunit activity. The increase in magnitude is gradual, however, essentially proportional to the number of units contributing to the multiunit response. For instance, if the pair-wise correlations are 0.01, multiunit clusters consisting of 10–20 units would be needed to obtain r_{SC} values in the typical range (0.1–0.2, indicated by horizontal dashed lines).



It is important to note that excessively restrictive criteria in spike sorting can lead to an underestimation of r_{SC} . Sorting spike waveforms in extracellular recordings is essentially a decision about when noisy voltage traces are similar enough that they are likely to have come from a single neuron. We considered a simple scenario in which the waveforms from two neurons were recorded on separate electrodes and clearly distinct, but each waveform was corrupted by noise (Fig. 5a,b). We simulated increasingly stringent spike sorting by discarding a proportion of waveforms from each neuron. In this simulation, no spikes are mistakenly assigned to the other unit; changing the threshold simply alters the proportion of events that are accepted as valid spikes. As fewer spikes are accepted (meaning that the criterion for acceptance becomes more stringent), the measured value of r_{SC} decreases (Fig. 5c). Such oversorting (discarding valid waveforms) decreases r_{SC} because whether a spike is accepted as valid or not is a random event, dependent strictly on noise, and variability that is independent for the two cells weakens measured correlation.

The relationship between the measured r_{SC} ($r_{SC-oversort}$) and the proportion of spikes discarded is given by

$$r_{SC-oversort} = \frac{r_{SC-original}}{\sqrt{\left(1 + \frac{p_1 \langle n_1 \rangle}{(1-p_1) \text{var}(n_1)}\right) \left(1 + \frac{p_2 \langle n_2 \rangle}{(1-p_2) \text{var}(n_2)}\right)}} \quad (2)$$

where p_1 and p_2 are the probabilities that a spike is discarded from neurons 1 and 2, respectively, and n_1 and n_2 are the spike counts of those cells (see **Supplementary Results** for derivation). In our simulation, the probability of deletion, p , is the same for the two cells ($p_1 = p_2$), the two cells have equal rates and the Fano factor equals 1, so $\langle n_1 \rangle = \langle n_2 \rangle = \text{var}(n_1) = \text{var}(n_2)$. In this case,

$$r_{SC-oversort} = (1-p)r_{SC-original} \quad (3)$$

Thus, when half the spikes are accepted, measured r_{SC} decreases twofold (from 0.20 to 0.10). When the probability of discarding waveforms is different for the two cells ($p_1 \neq p_2$) or when the Fano factors are different from 1, the decay with sort stringency will be different from that depicted in **Figure 5a–c**. Because the denominator in equation (2) is always greater than 1, however, $r_{SC-oversort}$ will always underestimate $r_{SC-original}$.

A second scenario in which spike sorting can reduce measured correlations is when waveforms belonging to a single neuron are mistakenly assigned to multiple neurons (Fig. 5d–f). We simulated this scenario by randomly dividing the spikes from a single neuron into two units (Fig. 5d). We then assessed the correlations between these units and a single unit (Fig. 5e) measured on another electrode. This manipulation reduced r_{SC} by roughly 30% compared with the true underlying correlation (Fig. 5f). When more than two units are created from a single unit, the measured value of r_{SC} falls further (Fig. 5f).

Variability in internal states can affect r_{SC} . Measurements of r_{SC} are based on sets of trials in which the stimulus and behavioral conditions are held as constant as possible. Despite experimenters’ best efforts, however, internal factors, such as arousal, attention or motivation, are bound to vary. Similarly, in experiments using anesthetics, the depth of anesthesia may vary over time. Such fluctuations could, in principle, co-modulate the responses of groups of cells and thus directly contribute to measurements of r_{SC} .

It is impossible to experimentally determine the degree to which this is the case, as internal variables, by definition, are not under experimental control. However, comparing the timescale of fluctuations in internal states with that of correlations can provide important constraints on their contribution. In the absence of salient changes in the visual scene, animals can only shift their attention approximately once every 400 ms^{43–45}. Even shifts in exogenous attention take 100–200 ms following an abrupt stimulus change^{43,44,46–49}. Fluctuations in other cognitive states, such as arousal or motivation, or in anesthetic state likely occur even more slowly.

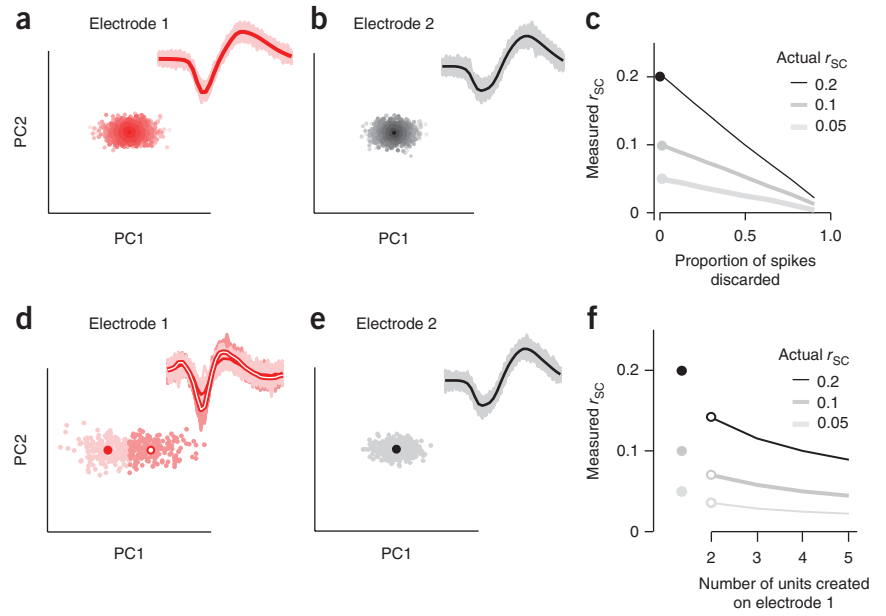
In contrast, correlations are dominated by fluctuations on shorter timescales. The timescale of correlation can be estimated in a trial by comparing r_{SC} computed in time windows of different sizes or by using the spike train crosscorrelogram^{8,12,22}. The contribution of fluctuations that occur across trials can be estimated by calculating r_{SC} between responses measured on different but nearby trials (for example, by shifting the trials of one of the two neurons)^{22,28}. These measures reveal that correlations are typically dominated by fluctuations on the timescale of tens to a few hundred milliseconds; correlations between responses measured in different trials are usually near zero^{22,28} (A.K., M.A. Smith, M.R.C. and J.H.R. Maunsell, unpublished data). Thus, although variations in cognitive factors or anesthesia states may well contribute to measured correlations, these correlations likely arise in large part from fluctuations on faster timescales.

Reconciling measurements of r_{SC}

The majority of studies that have measured spike count correlations have recorded neurons in primate sensory cortex. These studies report mean r_{SC} values in the range of 0.1–0.2 in pairs of similarly tuned, well-driven, nearby neurons (**Table 1**). In an apparent exception, a recent study by Ecker *et al.*²⁸ found mean correlations in primary visual cortex (V1) that were, on average, slightly positive but very close to zero. The authors concluded that variability in cortical neurons is independent and that the nonzero r_{SC} measurements in many other studies were the results of a variety of experimental artifacts.

The authors offer two primary explanations for this discrepancy. First, they argue that measured correlations in other studies were

Figure 5 Spike sorting errors can reduce the strength of measured correlations. (a,b) Two ensembles of spike waveforms (red and gray) created by taking two differently shaped waveforms (thick lines) and corrupting them with multiplicative and additive noise. These waveforms are represented by the amplitude of their first two principal components (PCs); each dot indicates one waveform. Lighter shading indicates greater distance in principal component space from the average waveform. (c) Effect of increasingly stringent sorting thresholds (that is, keeping only those spikes that are in a certain distance, in principal component space, of the average waveform) on the measured correlation strength, for true correlation values of 0.05, 0.1 and 0.2. As the proportion of discarded spikes increases, the measured value of r_{SC} decreases strongly. (d,e) Data are presented as in a and b for waveforms that form a tight cluster (e) and for those with a wider range of shapes (d). (f) Measured correlation (r_{SC}) between the well-isolated unit (e) and multiple random divisions of the waveforms in d. The true correlation is indicated by the filled dots to the left (0.05, 0.1 and 0.2). When the waveforms of the cell shown in d are assigned arbitrarily to two units (filled symbols), measured r_{SC} decreases by roughly 25%, compared with the true correlation between the pair. When the waveforms are arbitrarily divided among more than two units, r_{SC} falls further. These scenarios correspond to setting $p_1 = 0$ and $p_2 = 0.5$ (for two units created), 0.66 (for three units created), and so on, in equation (2).



inflated because of poor spike sorting, either by mistakenly swapping spikes between two neurons or by misidentifying multiunits as single units. The former error is only possible when correlations are measured between pairs of neurons recorded on the same electrode, the situation considered in their spike sorting simulations. Of the studies shown in **Table 1**, only one relies on pairs recorded in this manner⁶, so this cannot explain the departure from most previous results. That higher correlations in previous studies were the result of clustering single units into multiunits is also untenable. Previous studies would have needed to have mistakenly grouped responses from roughly 10–20 units with the near-zero correlations reported by Ecker *et al.* to explain the difference in mean r_{SC} values (**Fig. 4**). It seems inconceivable that all previous studies made mistakes of this magnitude. Furthermore, paired intracellular recordings, for which spike sorting is not an issue, have shown significant membrane potential and spiking correlations^{16,35,36}.

Second, Ecker *et al.* suggest that the correlations in previous studies arise from small fluctuations in uncontrolled cognitive factors, such as the animal’s internal state or behavioral factors such as fixational eye movements, which produce ‘artifactual’ correlations. As discussed above, the timescale of correlations makes it unlikely that they arise primarily from fluctuations in internal states. Furthermore, unlike other studies that asked subjects to perform difficult behavioral tasks (providing some means of controlling and assessing cognitive state), the subjects in Ecker *et al.* performed an easy fixation task. Thus, the internal states of their subjects likely varied as much as, or more than, in previous studies. Fixational eye movements are also unlikely to be the primary source of measured correlations because positive correlations are reported in studies that specifically remove trials containing detectable eye movements^{8,15,50} and in those that use anesthetized, paralyzed animals^{12,26}. Fixational eye movements would also be expected to produce anti-correlations in some pairs (for example, cells with offset spatial receptive fields, so that an eye movement that reduces drive to one cell would increase it to the other), but correlations are typically near zero or slightly positive in such cases²⁶.

Why then are the data of Ecker *et al.* so different from previous findings? A notable feature of the data in this study is that the firing rates were unusually low. Precise values were not provided, but summary plots suggest that the mean firing rate was a few spikes per second and that many cells had rates as low as 0.1 spikes per s. Thus, neurons typically fired either a single spike, or no spikes, on each trial. As our simulations show, such weak responses, whether because neurons were only weakly driven or because of overly stringent spike sorting criteria, result in low measured values of r_{SC} . A second contribution may be the location of the recordings. Correlations in the input layers of V1 are weaker than in other layers (M.A. Smith and A.K., unpublished data; J. Hansen and V. Dragoi (University of Texas Houston), personal communication). If the data were recorded in part from those layers⁵¹, correlations would be weaker than in studies targeting other layers or visual areas.

Where do we go from here?

Comparing correlations across studies, brain areas and task conditions. The experimental factors discussed here, in addition to the influences of distance, tuning similarity and architecture, can explain much of the diversity of reported r_{SC} values. For instance, for studies in the primate visual system, differences in mean rate can account for more than 41% of the across-study variance in reported values of r_{SC} (**Fig. 6**). It is clear that these factors must be considered when comparing results across studies, cortical areas, and stimulus or behavioral conditions, although some are more likely to affect conclusions than others. For example, when conclusions are based on the relative magnitude of correlations across conditions in a single experiment, spike sorting errors are unlikely to have a critical role. The dependence of r_{SC} on firing rates, on the other hand, can be an important issue. Behavioral and stimulus manipulations often affect firing rates, and neurons in different areas or circuits may vary in their responsiveness to experimental manipulations. Alternative, rate-independent metrics can be useful when rates vary across conditions^{38,52}, but determining the effect of stimulus and behavioral manipulations on correlations is




Figure 6 Differences in mean rates duration predict much of the variability in r_{sc} measurements across studies. The figure plots average r_{sc} as a function of mean firing rate and measurement duration for 13 studies of correlations in primate visual cortex, for which these data were either provided or could be estimated from summary figures. When a range of values was reported (for example, from different stimulus or behavioral conditions or for pairs of neurons separated by different distances), we plotted either the average or most common value, and in some cases we estimated values from summary plots. The mean firing rate accounts for 41% of the cross-study variance in mean r_{sc} values. Differences in spike sorting were not included in this meta-analysis because sort quality is rarely quantified or discussed.

most straightforward in cases when firing rates do not change^{12,15} or when correlations decrease while firing rates increase⁷.

Studying the influence of correlations on behavior. To date, most studies of correlations have focused on their effect on population encoding and decoding, used their properties to make inferences about network architecture or made comparisons across conditions to study neuronal correlates of perceptual or cognitive events. Ultimately, the goal of studying correlations should also be to understand how they affect computations in cortex and directly contribute to behavior. The effort to establish a link between correlated variability in populations of sensory neurons, the activity of downstream neurons and perceptual decisions is in its early stages. However, the studies reviewed below suggest that correlated variability directly contributes to perception and provide promising avenues for future research.

The best evidence that correlated variability in sensory neurons affects behavior is the well-established observation that fluctuations in the responses of individual neurons are predictive of an animal's perceptual decision (choice probability)^{53–55}. Choice probability has been observed in a range of tasks and cortical areas, as early as primary visual cortex⁵⁶ (I. Kang and J.H.R. Maunsell (Harvard University), personal communication; H. Nienborg and B.G. Cumming (US National Institutes of Health), personal communication). If the activity of sensory neurons were independent, the responses of any individual neuron would be measurably correlated with behavior only if very few neurons contributed to perceptual decisions, making the probability of encountering these cells very low. A more plausible explanation is that because the responses of neurons are correlated, each neuron's response reflects the shared activity of a large group of cells that contribute to perceptual decisions.

Fluctuations in nonspike-based measures of brain activity such as scalp recordings^{57,58} and the blood oxygen level-dependent signal measured in functional magnetic resonance imaging studies^{59–64} are also predictive of behavior. Fluctuations in such coarse signals presumably reflect the summed activity of many cells, suggesting that they are driven by correlated variability. Their relationship to decisions suggests that the underlying neuronal fluctuations affect behavior.

The relationship between behavior and fluctuations in neural activity suggests that correlated variability in sensory neurons affects perceptual decisions. The correlated variability underlying choice probability was originally thought to reflect noise in a bottom-up (that is, sensory) signal, but recent evidence suggests that choice probability at least partially reflects top-down influences⁶⁵. This raises the possibility that correlations in sensory neurons reflect rather than drive decisions, although it seems more plausible that the purpose of top-down signals is to influence decisions through their effect on pools of sensory neurons. Examining the time course of choice probability reveals that sensory neurons are correlated with decisions, and therefore with each other, during nearly the entire response to the stimulus^{65–68}. Furthermore, both experimental^{69,70} and theoretical studies^{71–74} indicate that sensory information is integrated until just before a decision is reported. Together, these results suggest that the correlated variability is present during, and not solely after, the period when perceptual decisions are made. However, establishing a causal link between correlated variability and perceptual decisions will require further research and may require the development and application of techniques that allow experimenters to manipulate correlated activity directly.

Principled statistical models. Pearson's correlation is a simple, descriptive statistic. Its simplicity and the fact that it can be estimated from experimentally tractable amounts of data has made it the metric of choice for both experimentalists and theoreticians for exploring issues of population coding and network function and connectivity. However, it is important to note that pair-wise correlations provide a limited view of complex population responses.

Recent theoretical work has begun to offer tools for providing a more complete and statistically principled approach to measuring interactions among cells. Multivariate point process models provide a full description of joint response distributions, incorporating information about receptive field properties, spiking history and interactions among cells^{75–78}. These methods provide means of simulating population responses, assessing the relative importance of single-cell properties and network influences on measured responses, and performing model-based decoding. Models of this type have been used to show that many of the observed fluctuations in single retinal ganglion cells⁷⁹ or neurons in sensorimotor cortex⁸⁰ can be accounted for by the current and past responses of other cells in the network. Future applications of these methods will undoubtedly provide important insights into issues of population coding and network structure.

Concluding remarks

Despite the difficulties in measuring and interpreting correlations, we have made important strides in characterizing their properties and how these depend on network state, behavioral experience and cognitive state⁸¹. Understanding which properties of correlations are most important—for instance, the most relevant timescales and sets of neurons—will ultimately depend on determining how population responses are interpreted by downstream networks, how they affect the computations performed there and how they guide behavioral decisions. With recent

improvements in experimental techniques and computational methods, this understanding feels very much within reach.

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We are grateful to P. Latham for analytical descriptions for our simulations and for the derivation of the relationship between measurement window and correlations that are not based on common spikes and to R. Coen-Cagli for the derivation relating correlation to the proportion of spikes discarded during spike sorting. We thank R. Coen-Cagli, B. Cumming, M. Histed, X. Jia, K. Josic, J. Maunsell, A. Ni, H. Nienborg, A. Pouget, O. Schwartz, S. Tanabe, and J.A. Movshon and E. Simoncelli and members of their laboratories for helpful discussions and comments on an earlier version of the manuscript. This work was supported by US National Institutes of Health grants R01 EY016774 (A.K.) and K99 EY020844-01 (M.R.C.).

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

Published online at <http://www.nature.com/natureneuroscience/>.

Reprints and permissions information is available online at <http://www.nature.com/reprints/index.html>.

1. Tolhurst, D.J., Movshon, J.A. & Dean, A.F. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* **23**, 775–785 (1983).
2. Shadlen, M.N. & Newsome, W.T. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* **18**, 3870–3896 (1998).
3. Abbott, L.F. & Dayan, P. The effect of correlated variability on the accuracy of a population code. *Neural Comput.* **11**, 91–101 (1999).
4. Averbeck, B.B., Latham, P.E. & Pouget, A. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, 358–366 (2006).
5. Nirenberg, S. & Latham, P.E. Decoding neuronal spike trains: how important are correlations? *Proc. Natl. Acad. Sci. USA* **100**, 7348–7353 (2003).
6. Zohary, E., Shadlen, M.N. & Newsome, W.T. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* **370**, 140–143 (1994).
7. Cohen, M.R. & Maunsell, J.H. Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* **12**, 1594–1600 (2009).
8. Mitchell, J.F., Sundberg, K.A. & Reynolds, J.H. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* **63**, 879–888 (2009).
9. Aertsen, A.M., Gerstein, G.L., Habib, M.K. & Palm, G. Dynamics of neuronal firing correlation: modulation of “effective connectivity”. *J. Neurophysiol.* **61**, 900–917 (1989).
10. Ahissar, E. *et al.* Dependence of cortical plasticity on correlated activity of single neurons and on behavioral context. *Science* **257**, 1412–1415 (1992).
11. Espinosa, I.E. & Gerstein, G.L. Cortical auditory neuron interactions during presentation of 3-tone sequences: effective connectivity. *Brain Res.* **450**, 39–50 (1988).
12. Kohn, A. & Smith, M.A. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* **25**, 3661–3673 (2005).
13. Gutnisky, D.A. & Dragoi, V. Adaptive coding of visual information in neural populations. *Nature* **452**, 220–224 (2008).
14. Komiyama, T. *et al.* Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. *Nature* **464**, 1182–1186 (2010).
15. Cohen, M.R. & Newsome, W.T. Context-dependent changes in functional circuitry in visual area MT. *Neuron* **60**, 162–173 (2008).
16. Poulet, J.F. & Petersen, C.C. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* **454**, 881–885 (2008).
17. Vaadia, E. *et al.* Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* **373**, 515–518 (1995).
18. Serié, P., Latham, P.E. & Pouget, A. Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. *Nat. Neurosci.* **7**, 1129–1135 (2004).
19. Greschner, M. *et al.* Correlated firing among major ganglion cell types in primate retina. *J. Physiol. (Lond.)* **589**, 75–86 (2011).
20. Reid, R.C. & Alonso, J.M. Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* **378**, 281–284 (1995).
21. Alonso, J.M. & Martinez, L.M. Functional connectivity between simple cells and complex cells in cat striate cortex. *Nat. Neurosci.* **1**, 395–403 (1998).
22. Bair, W., Zohary, E. & Newsome, W.T. Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J. Neurosci.* **21**, 1676–1697 (2001).
23. Reich, D.S., Mechler, F. & Victor, J.D. Independent and redundant information in nearby cortical neurons. *Science* **294**, 2566–2568 (2001).
24. Constantinidis, C. & Goldman-Rakic, P.S. Correlated discharges among putative pyramidal neurons and interneurons in the primate prefrontal cortex. *J. Neurophysiol.* **88**, 3487–3497 (2002).

25. Lee, D., Port, N.L., Kruse, W. & Georgopoulos, A.P. Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J. Neurosci.* **18**, 1161–1170 (1998).
26. Smith, M.A. & Kohn, A. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* **28**, 12591–12603 (2008).
27. Averbeck, B.B. & Lee, D. Neural noise and movement-related codes in the macaque supplementary motor area. *J. Neurosci.* **23**, 7630–7641 (2003).
28. Ecker, A.S. *et al.* Decorrelated neuronal firing in cortical microcircuits. *Science* **327**, 584–587 (2010).
29. Huang, X. & Lisberger, S.G. Noise correlations in cortical area MT and their potential impact on trial-by-trial variation in the direction and speed of smooth-pursuit eye movements. *J. Neurophysiol.* **101**, 3012–3030 (2009).
30. Jermakowicz, W.J., Chen, X., Khaytin, I., Bonds, A.B. & Casagrande, V.A. Relationship between spontaneous and evoked spike-time correlations in primate visual cortex. *J. Neurophysiol.* **101**, 2279–2289 (2009).
31. Rasch, M.J., Schuch, K., Logothetis, N.K. & Maass, W. Statistical comparison of spike responses to natural stimuli in monkey area V1 with simulated responses of a detailed laminar network model for a patch of V1. *J. Neurophysiol.* **105**, 757–778 (2011).
32. Zhang, M. & Alloway, K.D. Stimulus-induced intercolumnar synchronization of neuronal activity in rat barrel cortex: a laminar analysis. *J. Neurophysiol.* **92**, 1464–1478 (2004).
33. Renart, A. *et al.* The asynchronous state in cortical circuits. *Science* **327**, 587–590 (2010).
34. de la Rocha, J., Doiron, B., Shea-Brown, E., Josic, K. & Reyes, A. Correlation between neural spike trains increases with firing rate. *Nature* **448**, 802–806 (2007).
35. Lampl, I., Reichova, I. & Ferster, D. Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron* **22**, 361–374 (1999).
36. Okun, M. & Lampl, I. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat. Neurosci.* **11**, 535–537 (2008).
37. Kazama, H. & Wilson, R.I. Origins of correlated activity in an olfactory circuit. *Nat. Neurosci.* **12**, 1136–1144 (2009).
38. Dorn, J.D. & Ringach, D.L. Estimating membrane voltage correlations from extracellular spike trains. *J. Neurophysiol.* **89**, 2271–2278 (2003).
39. Carandini, M. Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biol.* **2**, e264 (2004).
40. Mazurek, M.E. & Shadlen, M.N. Limits to the temporal fidelity of cortical spike rate signals. *Nat. Neurosci.* **5**, 463–471 (2002).
41. Bedenbaugh, P. & Gerstein, G.L. Multiunit normalized cross correlation differs from the average single-unit normalized correlation. *Neural Comput.* **9**, 1265–1275 (1997).
42. Rosenbaum, R.J., Trousdale, J. & Josic, K. Pooling and correlated neural activity. *Front. Comput. Neurosci.* **4**, 9 (2010).
43. Müller, H.J. & Rabbitt, P.M. Reflexive and voluntary orienting of visual attention: time course of activation and resistance to interruption. *J. Exp. Psychol. Hum. Percept. Perform.* **15**, 315–330 (1989).
44. Cheal, M. & Lyon, D.R. Central and peripheral precuing of forced-choice discrimination. *Q. J. Exp. Psychol. A* **43**, 859–880 (1991).
45. Müller, M.M., Teder-Salejari, W. & Hillyard, S.A. The time course of cortical facilitation during cued shifts of spatial attention. *Nat. Neurosci.* **1**, 631–634 (1998).
46. Kröse, B.J. & Julesz, B. The control and speed of shifts of attention. *Vision Res.* **29**, 1607–1619 (1989).
47. Nakayama, K. & Mackeben, M. Sustained and transient components of focal visual attention. *Vision Res.* **29**, 1631–1647 (1989).
48. Bisley, J.W. & Goldberg, M.E. Neuronal activity in the lateral intraparietal area and spatial attention. *Science* **299**, 81–86 (2003).
49. Herrington, T.M. & Assad, J.A. Neural activity in the middle temporal area and lateral intraparietal area during endogenously cued shifts of attention. *J. Neurosci.* **29**, 14160–14176 (2009).
50. Bair, W. & O’Keefe, L.P. The influence of fixational eye movements on the response of neurons in area MT of the macaque. *Vis. Neurosci.* **15**, 779–786 (1998).
51. Berens, P., Keliris, G.A., Ecker, A.S., Logothetis, N.K. & Tolias, A.S. Feature selectivity of the gamma-band of the local field potential in primate primary visual cortex. *Front. Neurosci.* **2**, 199–207 (2008).
52. Amari, S. Measure of correlation orthogonal to change in firing rate. *Neural Comput.* **21**, 960–972 (2009).
53. Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S. & Movshon, J.A. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis. Neurosci.* **13**, 87–100 (1996).
54. Nienborg, H. & Cumming, B. Correlations between the activity of sensory neurons and behavior: how much do they tell us about a neuron’s causality? *Curr. Opin. Neurobiol.* **20**, 376–381 (2010).
55. Parker, A.J. & Newsome, W.T. Sense and the single neuron: probing the physiology of perception. *Annu. Rev. Neurosci.* **21**, 227–277 (1998).
56. Palmer, C., Cheng, S.Y. & Seidemann, E. Linking neuronal and behavioral performance in a reaction-time visual detection task. *J. Neurosci.* **27**, 8122–8137 (2007).
57. Bollimunta, A., Chen, Y., Schroeder, C.E. & Ding, M. Neuronal mechanisms of cortical alpha oscillations in awake-behaving macaques. *J. Neurosci.* **28**, 9976–9988 (2008).
58. Thut, G., Nietzel, A., Brandt, S.A. & Pascual-Leone, A. Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *J. Neurosci.* **26**, 9494–9502 (2006).



59. Fox, M.D., Snyder, A.Z., Vincent, J.L. & Raichle, M.E. Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron* **56**, 171–184 (2007).
60. Grill-Spector, K., Knouf, N. & Kanwisher, N. The fusiform face area subserves face perception, not generic within-category identification. *Nat. Neurosci.* **7**, 555–562 (2004).
61. Leber, A.B. Neural predictors of within-subject fluctuations in attentional control. *J. Neurosci.* **30**, 11458–11465 (2010).
62. Ress, D., Backus, B.T. & Heeger, D.J. Activity in primary visual cortex predicts performance in a visual detection task. *Nat. Neurosci.* **3**, 940–945 (2000).
63. Ress, D. & Heeger, D.J. Neuronal correlates of perception in early visual cortex. *Nat. Neurosci.* **6**, 414–420 (2003).
64. Sapiro, A., d'Avossa, G., McAvoy, M., Shulman, G.L. & Corbetta, M. Brain signals for spatial attention predict performance in a motion discrimination task. *Proc. Natl. Acad. Sci. USA* **102**, 17810–17815 (2005).
65. Nienborg, H. & Cumming, B.G. Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* **459**, 89–92 (2009).
66. Cohen, M.R. & Newsome, W.T. Estimates of the contribution of single neurons to perception depend on timescale and noise correlation. *J. Neurosci.* **29**, 6635–6648 (2009).
67. Cook, E.P. & Maunsell, J.H. Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nat. Neurosci.* **5**, 985–994 (2002).
68. Price, N.S. & Born, R.T. Timescales of sensory- and decision-related activity in the middle temporal and medial superior temporal areas. *J. Neurosci.* **30**, 14036–14045 (2010).
69. Ditterich, J., Mazurek, M.E. & Shadlen, M.N. Microstimulation of visual cortex affects the speed of perceptual decisions. *Nat. Neurosci.* **6**, 891–898 (2003).
70. Huk, A.C. & Shadlen, M.N. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J. Neurosci.* **25**, 10420–10436 (2005).
71. Beck, J.M. *et al.* Probabilistic population codes for Bayesian decision making. *Neuron* **60**, 1142–1152 (2008).
72. Mazurek, M.E., Roitman, J.D., Ditterich, J. & Shadlen, M.N. A role for neural integrators in perceptual decision making. *Cereb. Cortex* **13**, 1257–1269 (2003).
73. Wang, X.J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* **36**, 955–968 (2002).
74. Wong, K.F., Huk, A.C., Shadlen, M.N. & Wang, X.J. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front. Comput. Neurosci.* **1**, 6 (2007).
75. Truccolo, W., Eden, U.T., Fellows, M.R., Donoghue, J.P. & Brown, E.N. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J. Neurophysiol.* **93**, 1074–1089 (2005).
76. Kass, R.E., Ventura, V. & Brown, E.N. Statistical issues in the analysis of neuronal data. *J. Neurophysiol.* **94**, 8–25 (2005).
77. Paninski, L. *et al.* A new look at state-space models for neural data. *J. Comput. Neurosci.* **29**, 107–126 (2010).
78. Okatan, M., Wilson, M.A. & Brown, E.N. Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Comput.* **17**, 1927–1961 (2005).
79. Pillow, J.W. *et al.* Spatio-temporal correlations and visual signaling in a complete neuronal population. *Nature* **454**, 995–999 (2008).
80. Truccolo, W., Hochberg, L.R. & Donoghue, J.P. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nat. Neurosci.* **13**, 105–111 (2010).
81. Kohn, A., Zandvakili, A. & Smith, M.A. Correlations and brain states: from electrophysiology to functional imaging. *Curr. Opin. Neurobiol.* **19**, 434–438 (2009).
82. Poort, J. & Roelfsema, P.R. Noise correlations have little influence on the coding of selective attention in area V1. *Cereb. Cortex* **19**, 543–553 (2009).
83. Samonds, J.M., Potetz, B.R. & Lee, T.S. Cooperative and competitive interactions facilitate stereo computations in macaque primary visual cortex. *J. Neurosci.* **29**, 15780–15795 (2009).
84. Erickson, C.A., Jagadeesh, B. & Desimone, R. Clustering of perirhinal neurons with similar properties following visual experience in adult monkeys. *Nat. Neurosci.* **3**, 1143–1148 (2000).
85. Averbeck, B.B. & Lee, D. Effects of noise correlations on information encoding and decoding. *J. Neurophysiol.* **95**, 3633–3644 (2006).
86. Stark, E., Globerson, A., Asher, I. & Abeles, M. Correlations between groups of premotor neurons carry information about prehension. *J. Neurosci.* **28**, 10618–10630 (2008).
87. Maynard, E.M. *et al.* Neuronal interactions improve cortical population coding of movement direction. *J. Neurosci.* **19**, 8083–8093 (1999).
88. Nevet, A., Morris, G., Saban, G., Arkadir, D. & Bergman, H. Lack of spike-count and spike-time correlations in the substantia nigra reticulata despite overlap of neural responses. *J. Neurophysiol.* **98**, 2232–2243 (2007).
89. Cohen, J.Y. *et al.* Cooperation and competition among frontal eye field neurons during visual target selection. *J. Neurosci.* **30**, 3227–3238 (2010).
90. Bichot, N.P., Thompson, K.G., Chenchal Rao, S. & Schall, J.D. Reliability of macaque frontal eye field neurons signaling saccade targets during visual search. *J. Neurosci.* **21**, 713–725 (2001).
91. Barlow, H.B. & Foldiak, P. Adaptation and decorrelation in the cortex. in *The Computing Neuron* (eds. Durbin, R., Miall, C. & Mitchinson, G.) (Addison-Wesley, New York, 1989).
92. Vinje, W.E. & Gallant, J.L. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* **287**, 1273–1276 (2000).