

Published in final edited form as:

Nat Neurosci. ; 14(7): 811–819. doi:10.1038/nn.2842.

Measuring and interpreting neuronal correlations

Marlene R. Cohen¹ and Adam Kohn²

¹ Department of Neurobiology, Harvard Medical School, Boston MA

² Department of Neuroscience, Albert Einstein College of Medicine, Bronx NY

Abstract

Mounting evidence suggests that understanding how the brain encodes information and performs computations will require studying 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 to detailed exploration of their properties and contributions to cortical processing. Studies to date, however, 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 dramatically 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 impact 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 impact 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, since shared fluctuations in response cannot be averaged away⁶. More recently, it has been shown that attentional modulation of correlations accounts for more of its impact on sensory coding than modulation of firing rate^{7,8}. Determining how correlations are affected by stimulus drive⁹⁻¹², by learning or experience^{10,13,14}, or by changes in behavioral context^{7,8,15-17} is thus 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 within the retina¹⁹, between the visual thalamus and cortex²⁰, and between neurons within 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 instance, 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. 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 Figure 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 their spike count responses to repeated presentations of identical stimuli under the same behavioral conditions (spike count correlation or r_{SC} , also called noise correlation; Figure 1A,B).

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 (i.e. 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 is 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.

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 cortex (³²; M.A. Smith and A. Kohn, personal communication; J. Hansen and V. Dragoi, 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 strength²⁻⁶. 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 both for elucidating their role in sensory processing and for making inferences about the circuitry and mechanisms that generate them. For this reason, we explore below 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 under study^{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 where stimulus preferences are less understood. The time window used to count spikes is purely an experimental decision, and thus 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 cofluctuations 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}.

This relationship is illustrated with a simple simulation shown in Figure 2 (similar to³⁴). 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 non-linearity shown in the middle panel. 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 (Figure 2A), the correlation in spiking responses is similar to that of the subthreshold response. However, when the mean membrane potential is far below threshold (Figure 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 thus be more correlated (Figure 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 sp/s and the other 100 sp/s has the same geometric mean response as a pair whose mean responses are both 1 sp/s, but the measured r_{SC} of 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, Figure 2D has vertical and horizontal bands, 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 (Figure 3). To do so, it was necessary to use a framework in which we could control the timescale of correlation. Since the simulations of 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 (Figure 3A,B; see also⁴⁰).

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 (Figure 3A), the resulting synchrony will be evident in spike count correlations based on responses measured over arbitrarily small windows (line marked 'no jitter' in Figure 3C). If the times of these common spikes are jittered (Figure 3B), however, the resulting spike count correlation will only be fully evident when it is calculated from responses during longer response epochs (lines labeled 5-80 ms in Figure 3C). For instance, when spike times are jittered using a Gaussian distribution with a standard deviation 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 (see²² and Supplementary Material for analytical description). Note that in this scenario correlations do not arise from nearly synchronous spikes, yet 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.

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, while 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 to 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 pairwise correlations ranged from very small (0.001) to more typical (0.1) values (Figure 4).

These simulations show that when pairwise 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 pairwise 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 (i.e. as $nr_{SC\text{-pair}}$). If, for example, the underlying pairwise 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.

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 (Figure 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 (Figure 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} ($r_{SC\text{-oversort}}$) and the proportion of spikes discarded is given by

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

where p_1, p_2 are the probabilities that a spike is discarded from neuron 1 and 2, respectively, and n_1, n_2 are the spike counts of those cells (see Supplementary Material 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,

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

Therefore, 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 (Figures 5D-F). We simulated this scenario by randomly dividing the spikes from a single neuron into two units (different shades of red in Figure 5D). We then assessed the correlations between these units and a single unit (Figure 5E) measured on another electrode. This manipulation reduced r_{SC} by roughly 30% compared to the true underlying correlation (Figure 5F, open symbols compared to filled). When more than two units are created from a single unit, the measured value of r_{SC} falls further (Figure 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 comodulate the responses of groups of cells and thus directly contribute to measurements of r_{SC} .

It is impossible to determine experimentally 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 ms⁴³⁻⁴⁵. Even shifts in exogenous attention take 100-200 ms following an abrupt stimulus change^{43, 44, 46-49}. Fluctuations in other cognitive states like 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 within 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 (e.g. 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}; Kohn and Smith, Cohen and Maunsell, unpublished observations). 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 found mean correlations in primary visual cortex (V1) that were on average slightly positive but very close to zero²⁸. The authors concluded that variability in cortical neurons is independent and that the non-zero r_{SC} measurements in many other studies were due to a variety of experimental artifacts.

The authors offer two primary explanations for this discrepancy. First, the authors argue that measured correlations in other studies were 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 in Table 1, only one relies on pairs recorded this way⁶, so this explanation 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 (Figure 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 like fixational eye movements, which produce “artificial” 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. Therefore, 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 (e.g. 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 striking 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 many cells had rates as low as 0.1 sp/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. Kohn, personal communication; J. Hansen and V. Dragoi, 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 33% of the across-study variance in reported values of r_{SC} (Figure 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 instance, when conclusions are based on the relative magnitude of

correlations across conditions in a single experiment, spike sorting errors are unlikely to play a critical role. The dependence of r_{SC} on firing rates, on the other hand, can be a key 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 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 impact 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 cortex^(56; I. Kang and J.H.R. Maunsell, personal communication; H.N. and B.G. Cumming, 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 non-spike based measures of brain activity such as scalp recordings^{57, 58} and the BOLD signal measured in fMRI studies⁵⁹⁻⁶⁴ 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 (i.e. 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 stimulus⁶⁵⁻⁶⁸. Furthermore, both experimental^{69, 70} and theoretical studies⁷¹⁻⁷⁴ 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 to explore issues of population coding and network function and connectivity. However, it is important to note that pairwise 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 cells⁷⁵⁻⁷⁸. 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.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We are grateful to Peter 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 Ruben Coen-Cagli for the derivation relating correlation to the proportion of spikes discarded during spike sorting. We thank Ruben Coen-Cagli, Bruce Cumming, Mark Histed, Xiaoxuan Jia, Kresimir Josic, John Maunsell, Amy Ni, Hendrikje Nienborg, Alex Pouget, Odelia Schwartz, Seiji Tanabe, as well as Tony Movshon and Eero Simoncelli and members of their laboratories for helpful discussions and comments on an earlier version of the manuscript. This work was supported by NIH grants R01 EY016774 (AK) and K99 EY020844-01 (MRC).

References

1. Tolhurst DJ, Movshon JA, Dean AF. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* 1983; 23:775–785. [PubMed: 6623937]
2. Shadlen MN, Newsome WT. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci.* 1998; 18:3870–3896. [PubMed: 9570816]

3. Abbott LF, Dayan P. The effect of correlated variability on the accuracy of a population code. *Neural Comput.* 1999; 11:91–101. [PubMed: 9950724]
4. Averbeck BB, Latham PE, Pouget A. Neural correlations, population coding and computation. *Nat Rev Neurosci.* 2006; 7:358–366. [PubMed: 16760916]
5. Nirenberg S, Latham PE. Decoding neuronal spike trains: how important are correlations? *Proc Natl Acad Sci U S A.* 2003; 100:7348–7353. [PubMed: 12775756]
6. Zohary E, Shadlen MN, Newsome WT. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature.* 1994; 370:140–143. [PubMed: 8022482]
7. Cohen MR, Maunsell JH. Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci.* 2009; 12:1594–1600. [PubMed: 19915566]
8. Mitchell JF, Sundberg KA, Reynolds JH. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron.* 2009; 63:879–888. [PubMed: 19778515]
9. Aertsen AM, Gerstein GL, Habib MK, Palm G. Dynamics of neuronal firing correlation: modulation of “effective connectivity”. *J Neurophysiol.* 1989; 61:900–917. [PubMed: 2723733]
10. Ahissar E, et al. Dependence of cortical plasticity on correlated activity of single neurons and on behavioral context. *Science.* 1992; 257:1412–1415. [PubMed: 1529342]
11. Espinosa IE, Gerstein GL. Cortical auditory neuron interactions during presentation of 3-tone sequences: effective connectivity. *Brain Res.* 1988; 450:39–50. [PubMed: 3401720]
12. Kohn A, Smith MA. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci.* 2005; 25:3661–3673. [PubMed: 15814797]
13. Gutnisky DA, Dragoi V. Adaptive coding of visual information in neural populations. *Nature.* 2008; 452:220–224. [PubMed: 18337822]
14. Komiyama T, et al. Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. *Nature.* 2010; 464:1182–1186. [PubMed: 20376005]
15. Cohen MR, Newsome WT. Context-dependent changes in functional circuitry in visual area MT. *Neuron.* 2008; 60:162–173. [PubMed: 18940596]
16. Poulet JF, Petersen CC. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature.* 2008; 454:881–885. [PubMed: 18633351]
17. Vaadia E, et al. Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature.* 1995; 373:515–518. [PubMed: 7845462]
18. Series P, Latham PE, Pouget A. Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. *Nat Neurosci.* 2004; 7:1129–1135. [PubMed: 15452579]
19. Greschner M, et al. Correlated firing among major ganglion cell types in primate retina. *J Physiol.* 589:75–86. [PubMed: 20921200]
20. Reid RC, Alonso JM. Specificity of monosynaptic connections from thalamus to visual cortex. *Nature.* 1995; 378:281–284. [PubMed: 7477347]
21. Alonso JM, Martinez LM. Functional connectivity between simple cells and complex cells in cat striate cortex. *Nat Neurosci.* 1998; 1:395–403. [PubMed: 10196530]
22. Bair W, Zohary E, Newsome WT. Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J Neurosci.* 2001; 21:1676–1697. [PubMed: 11222658]
23. Reich DS, Mechler F, Victor JD. Independent and redundant information in nearby cortical neurons. *Science.* 2001; 294:2566–2568. [PubMed: 11752580]
24. Constantinidis C, Goldman-Rakic PS. Correlated discharges among putative pyramidal neurons and interneurons in the primate prefrontal cortex. *J Neurophysiol.* 2002; 88:3487–3497. [PubMed: 12466463]
25. Lee D, Port NL, Kruse W, Georgopoulos AP. Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J Neurosci.* 1998; 18:1161–1170. [PubMed: 9437036]
26. Smith MA, Kohn A. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J Neurosci.* 2008; 28:12591–12603. [PubMed: 19036953]
27. Averbeck BB, Lee D. Neural noise and movement-related codes in the macaque supplementary motor area. *J Neurosci.* 2003; 23:7630–7641. [PubMed: 12930802]

28. Ecker AS, et al. Decorrelated neuronal firing in cortical microcircuits. *Science*. 2010; 327:584–587. [PubMed: 20110506]
29. Huang X, Lisberger SG. 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*. 2009; 101:3012–3030. [PubMed: 19321645]
30. Jermakowicz WJ, Chen X, Khaytin I, Bonds AB, Casagrande VA. Relationship between spontaneous and evoked spike-time correlations in primate visual cortex. *J Neurophysiol*. 2009; 101:2279–2289. [PubMed: 19211656]
31. Rasch MJ, Schuch K, Logothetis NK, 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. [PubMed: 21106898]
32. Zhang M, Alloway KD. Stimulus-induced intercolumnar synchronization of neuronal activity in rat barrel cortex: a laminar analysis. *J Neurophysiol*. 2004; 92:1464–1478. [PubMed: 15056676]
33. Renart A, et al. The asynchronous state in cortical circuits. *Science*. 327:587–590. [PubMed: 20110507]
34. de la Rocha J, Doiron B, Shea-Brown E, Josic K, Reyes A. Correlation between neural spike trains increases with firing rate. *Nature*. 2007; 448:802–806. [PubMed: 17700699]
35. Lampl I, Reichova I, Ferster D. Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron*. 1999; 22:361–374. [PubMed: 10069341]
36. Okun M, Lampl I. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat Neurosci*. 2008; 11:535–537. [PubMed: 18376400]
37. Kazama H, Wilson RI. Origins of correlated activity in an olfactory circuit. *Nat Neurosci*. 2009; 12:1136–1144. [PubMed: 19684589]
38. Dorn JD, Ringach DL. Estimating membrane voltage correlations from extracellular spike trains. *J Neurophysiol*. 2003; 89:2271–2278. [PubMed: 12686584]
39. Carandini M. Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biol*. 2004; 2:E264. [PubMed: 15328535]
40. Mazurek ME, Shadlen MN. Limits to the temporal fidelity of cortical spike rate signals. *Nat Neurosci*. 2002; 5:463–471. [PubMed: 11976706]
41. Bedenbaugh P, Gerstein GL. Multiunit normalized cross correlation differs from the average single-unit normalized correlation. *Neural Comput*. 1997; 9:1265–1275. [PubMed: 9248062]
42. Rosenbaum RJ, Trousdale J, Josic K. Pooling and correlated neural activity. *Front Comput Neurosci*. 4:9. [PubMed: 20485451]
43. Muller HJ, Rabbitt PM. Reflexive and voluntary orienting of visual attention: time course of activation and resistance to interruption. *J Exp Psychol Hum Percept Perform*. 1989; 15:315–330. [PubMed: 2525601]
44. Cheal M, Lyon DR. Central and peripheral precuing of forced-choice discrimination. *Q J Exp Psychol A*. 1991; 43:859–880. [PubMed: 1775667]
45. Muller MM, Teder-Salejarvi W, Hillyard SA. The time course of cortical facilitation during cued shifts of spatial attention. *Nat Neurosci*. 1998; 1:631–634. [PubMed: 10196572]
46. Krose BJ, Julesz B. The control and speed of shifts of attention. *Vision Res*. 1989; 29:1607–1619. [PubMed: 2635484]
47. Nakayama K, Mackeben M. Sustained and transient components of focal visual attention. *Vision Res*. 1989; 29:1631–1647. [PubMed: 2635486]
48. Bisley JW, Goldberg ME. Neuronal activity in the lateral intraparietal area and spatial attention. *Science*. 2003; 299:81–86. [PubMed: 12511644]
49. Herrington TM, Assad JA. Neural activity in the middle temporal area and lateral intraparietal area during endogenously cued shifts of attention. *J Neurosci*. 2009; 29:14160–14176. [PubMed: 19906965]
50. Bair W, O'Keefe LP. The influence of fixational eye movements on the response of neurons in area MT of the macaque. *Vis Neurosci*. 1998; 15:779–786. [PubMed: 9682878]

51. Berens P, Keliris GA, Ecker AS, Logothetis NK, Tolias AS. Feature selectivity of the gamma-band of the local field potential in primate primary visual cortex. *Front Neurosci.* 2008; 2:199–207. [PubMed: 19225593]
52. Amari S. Measure of correlation orthogonal to change in firing rate. *Neural Comput.* 2009; 21:960–972. [PubMed: 19018699]
53. Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci.* 1996; 13:87–100. [PubMed: 8730992]
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.* 2010; 20:376–381. [PubMed: 20545019]
55. Parker AJ, Newsome WT. Sense and the single neuron: probing the physiology of perception. *Annu Rev Neurosci.* 1998; 21:227–277. [PubMed: 9530497]
56. Palmer C, Cheng SY, Seidemann E. Linking neuronal and behavioral performance in a reaction-time visual detection task. *J Neurosci.* 2007; 27:8122–8137. [PubMed: 17652603]
57. Bollimunta A, Chen Y, Schroeder CE, Ding M. Neuronal mechanisms of cortical alpha oscillations in awake-behaving macaques. *J Neurosci.* 2008; 28:9976–9988. [PubMed: 18829955]
58. Thut G, Nietzel A, Brandt SA, Pascual-Leone A. Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *J Neurosci.* 2006; 26:9494–9502. [PubMed: 16971533]
59. Fox MD, Snyder AZ, Vincent JL, Raichle ME. Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron.* 2007; 56:171–184. [PubMed: 17920023]
60. Grill-Spector K, Knouf N, Kanwisher N. The fusiform face area subserves face perception, not generic within-category identification. *Nat Neurosci.* 2004; 7:555–562. [PubMed: 15077112]
61. Leber AB. Neural predictors of within-subject fluctuations in attentional control. *J Neurosci.* 30:11458–11465. [PubMed: 20739567]
62. Ress D, Backus BT, Heeger DJ. Activity in primary visual cortex predicts performance in a visual detection task. *Nat Neurosci.* 2000; 3:940–945. [PubMed: 10966626]
63. Ress D, Heeger DJ. Neuronal correlates of perception in early visual cortex. *Nat Neurosci.* 2003; 6:414–420. [PubMed: 12627164]
64. Sapir A, d'Avossa G, McAvoy M, Shulman GL, Corbetta M. Brain signals for spatial attention predict performance in a motion discrimination task. *Proc Natl Acad Sci U S A.* 2005; 102:17810–17815. [PubMed: 16306268]
65. Nienborg H, Cumming BG. Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature.* 2009; 459:89–92. [PubMed: 19270683]
66. Cohen MR, Newsome WT. Estimates of the contribution of single neurons to perception depend on timescale and noise correlation. *J Neurosci.* 2009; 29:6635–6648. [PubMed: 19458234]
67. Cook EP, Maunsell JH. Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nat Neurosci.* 2002; 5:985–994. [PubMed: 12244324]
68. Price NS, Born RT. Timescales of sensory- and decision-related activity in the middle temporal and medial superior temporal areas. *J Neurosci.* 30:14036–14045. [PubMed: 20962225]
69. Ditterich J, Mazurek ME, Shadlen MN. Microstimulation of visual cortex affects the speed of perceptual decisions. *Nat Neurosci.* 2003; 6:891–898. [PubMed: 12858179]
70. Huk AC, Shadlen MN. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci.* 2005; 25:10420–10436. [PubMed: 16280581]
71. Beck JM, et al. Probabilistic population codes for Bayesian decision making. *Neuron.* 2008; 60:1142–1152. [PubMed: 19109917]
72. Mazurek ME, Roitman JD, Ditterich J, Shadlen MN. A role for neural integrators in perceptual decision making. *Cereb Cortex.* 2003; 13:1257–1269. [PubMed: 14576217]
73. Wang XJ. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron.* 2002; 36:955–968. [PubMed: 12467598]

74. Wong KF, Huk AC, Shadlen MN, Wang XJ. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front Comput Neurosci.* 2007; 1:6. [PubMed: 18946528]
75. Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J Neurophysiol.* 2005; 93:1074–1089. [PubMed: 15356183]
76. Kass RE, Ventura V, Brown EN. Statistical issues in the analysis of neuronal data. *J Neurophysiol.* 2005; 94:8–25. [PubMed: 15985692]
77. Paninski L, et al. A new look at state-space models for neural data. *J Comput Neurosci.* 29:107–126. [PubMed: 19649698]
78. Okatan M, Wilson MA, Brown EN. Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Comput.* 2005; 17:1927–1961. [PubMed: 15992486]
79. Pillow JW, et al. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature.* 2008; 454:995–999. [PubMed: 18650810]
80. Truccolo W, Hochberg LR, Donoghue JP. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nat Neurosci.* 13:105–111. [PubMed: 19966837]
81. Kohn A, Zandvakili A, Smith MA. Correlations and brain states: from electrophysiology to functional imaging. *Curr Opin Neurobiol.* 2009
82. Barlow, HB.; Foldiak, P. Adaptation and decorrelation in the cortex.. In: Durbin, R.; Miall, C.; Mitchinson, G., editors. *The computing neuron.* Addison-Wesley; New York: 1989.
83. Vinje WE, Gallant JL. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science.* 2000; 287:1273–1276. [PubMed: 10678835]
84. Poort J, Roelfsema PR. Noise correlations have little influence on the coding of selective attention in area V1. *Cereb Cortex.* 2009; 19:543–553. [PubMed: 18552357]
85. Samonds JM, Potetz BR, Lee TS. Cooperative and competitive interactions facilitate stereo computations in macaque primary visual cortex. *J Neurosci.* 2009; 29:15780–15795. [PubMed: 20016094]
86. Erickson CA, Jagadeesh B, Desimone R. Clustering of perirhinal neurons with similar properties following visual experience in adult monkeys. *Nat Neurosci.* 2000; 3:1143–1148. [PubMed: 11036272]
87. Averbeck BB, Lee D. Effects of noise correlations on information encoding and decoding. *J Neurophysiol.* 2006; 95:3633–3644. [PubMed: 16554512]
88. Stark E, Globerson A, Asher I, Abeles M. Correlations between groups of premotor neurons carry information about prehension. *J Neurosci.* 2008; 28:10618–10630. [PubMed: 18923038]
89. Maynard EM, et al. Neuronal interactions improve cortical population coding of movement direction. *J Neurosci.* 1999; 19:8083–8093. [PubMed: 10479708]
90. 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.* 2007; 98:2232–2243. [PubMed: 17699698]
91. Cohen JY, et al. Cooperation and competition among frontal eye field neurons during visual target selection. *J Neurosci.* 30:3227–3238. [PubMed: 20203182]
92. Bichot NP, Thompson KG, Chenthal Rao S, Schall JD. Reliability of macaque frontal eye field neurons signaling saccade targets during visual search. *J Neurosci.* 2001; 21:713–725. [PubMed: 11160450]
93. Khatri V, Bruno RM, Simons DJ. Stimulus-specific and stimulus-nonspecific firing synchrony and its modulation by sensory adaptation in the whisker-to-barrel pathway. *J Neurophysiol.* 2009; 101:2328–2338. [PubMed: 19279146]

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, for instance through adaptation⁸² or contextual modulation⁸³, 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-correlogram (CCG; ²²).

Synchrony measures the extent to which the timing of individual spikes is precisely aligned, typically on the timescale of one or a few milliseconds. It is typically quantified as a measure of a sharp peak in the CCG. 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 past. 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.

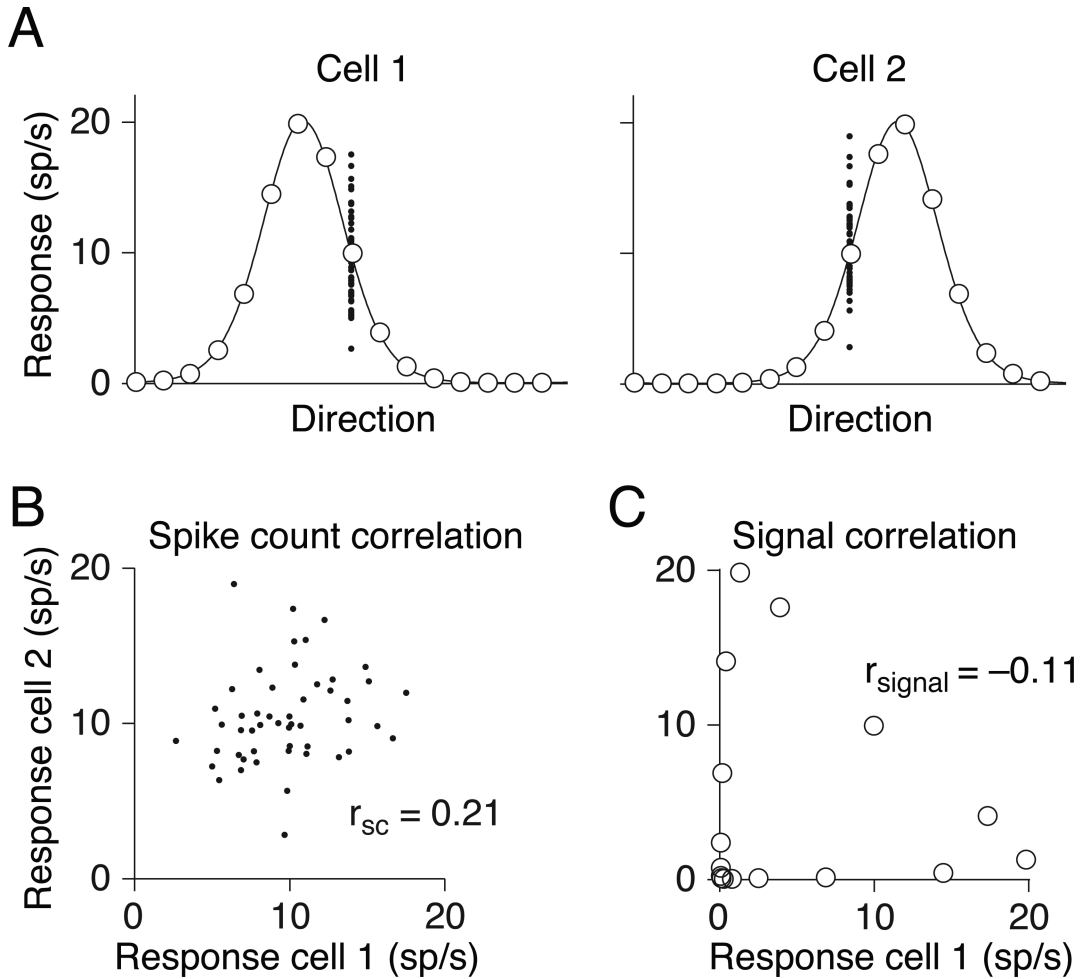


Figure 1.

Types of pairwise 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.

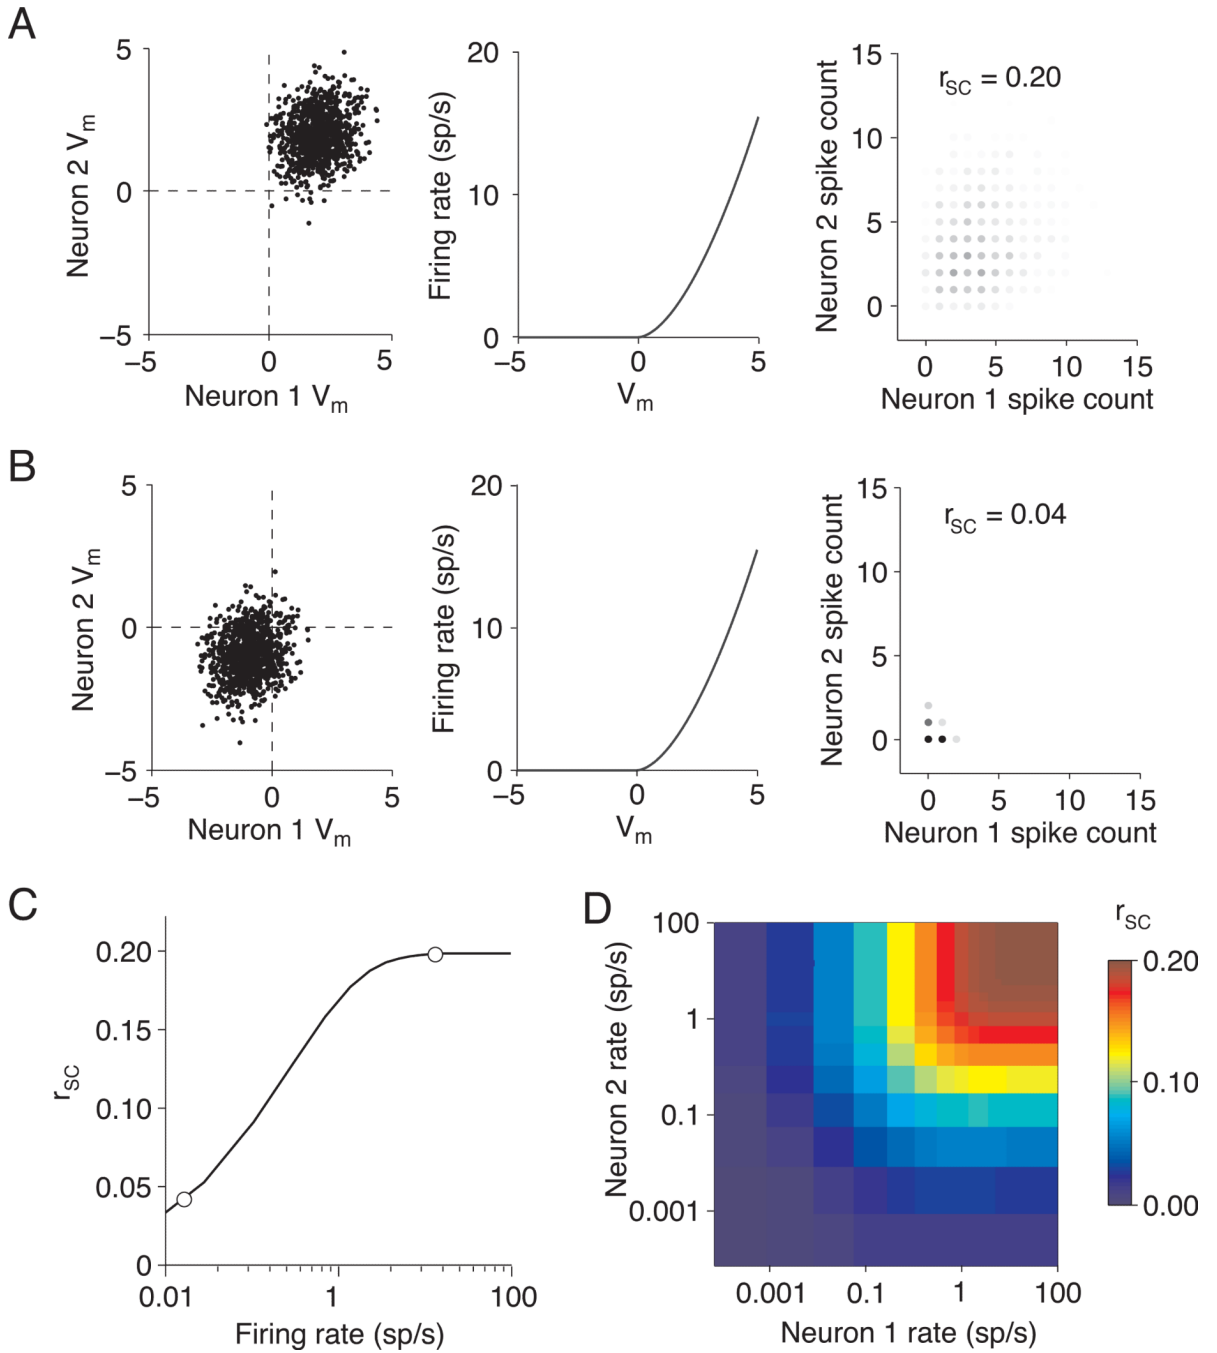


Figure 2.

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 panel). These voltage events were converted to extracellular firing rates by passing them through a non-linearity such that the firing rate on the i th trial $fr_i = Vm_i^{1.7}$ (center panel). 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 panel). In the case of high mean voltages, the measured spike count correlation is close to the input correlation

($r_{SC} = 0.20$). **B.** Same as 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 counts 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.

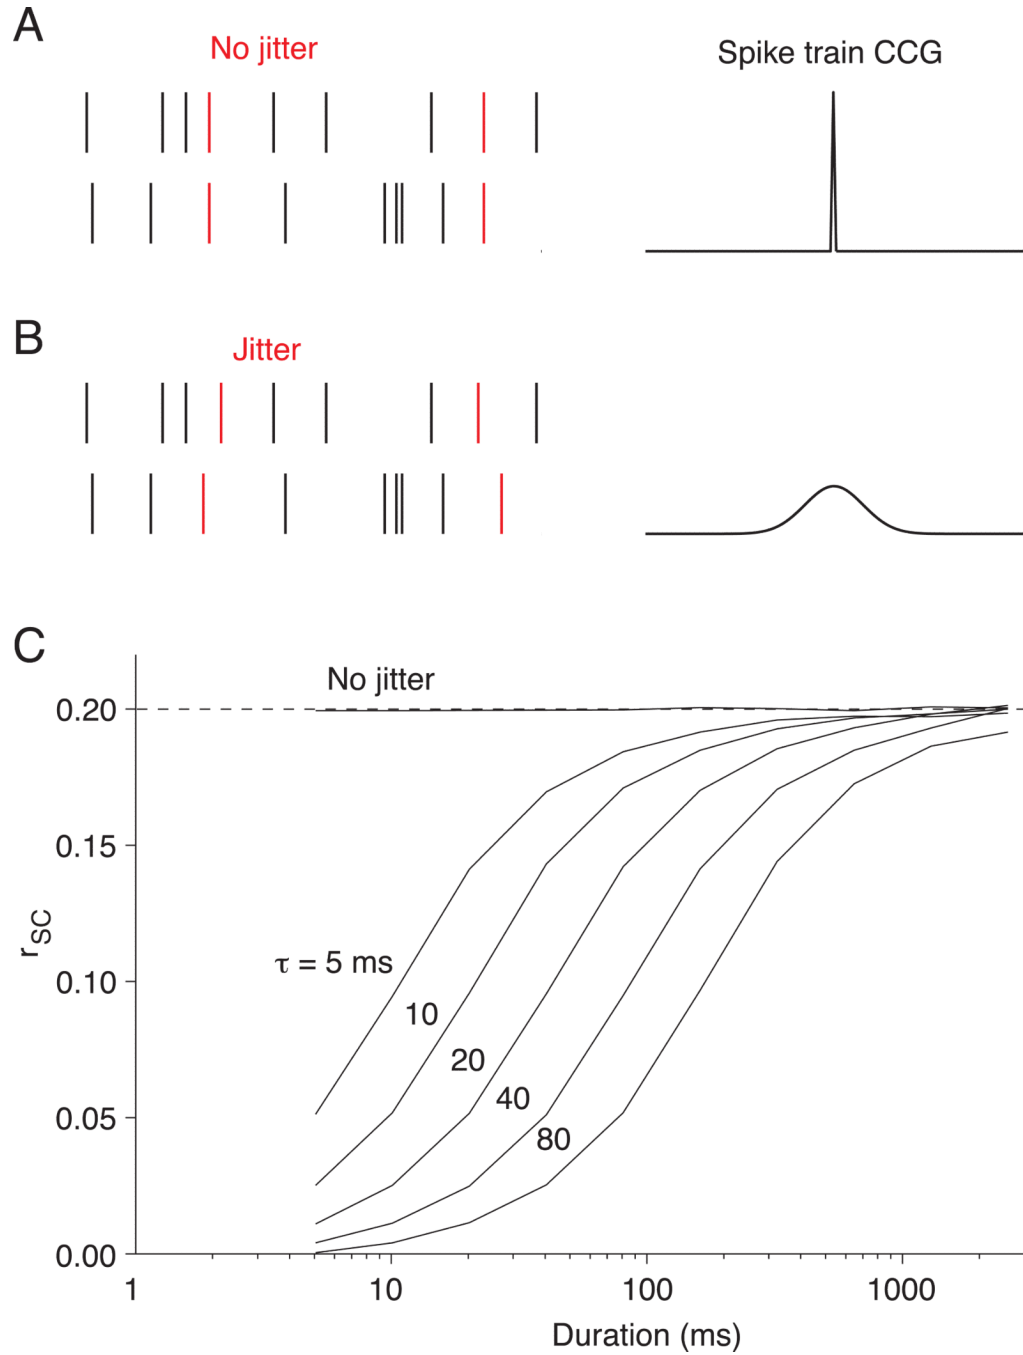


Figure 3.

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 panel; mean = 20 sp/s) and inserted shared, synchronous spikes (red spikes, mean = 5 sp/s). Therefore, each neuron's response was the sum of independent and shared spikes (mean = 25 sp/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 panel). **B.** Same as A, for longer jitter times. We simulated variable timescales by jittering the timing of the correlated spikes by an amount picked from a Gaussian distribution (left panel, red spikes).

This results in a cross correlogram with a peak whose width depends on the standard deviation of the Gaussian distribution (right panel). C. Measured r_{SC} as a function of counting window for several timescales of correlation. The number next to each curve corresponds to the standard deviation 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 pairwise correlation values of 0.001-0.3. 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 pairwise 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 dotted lines).

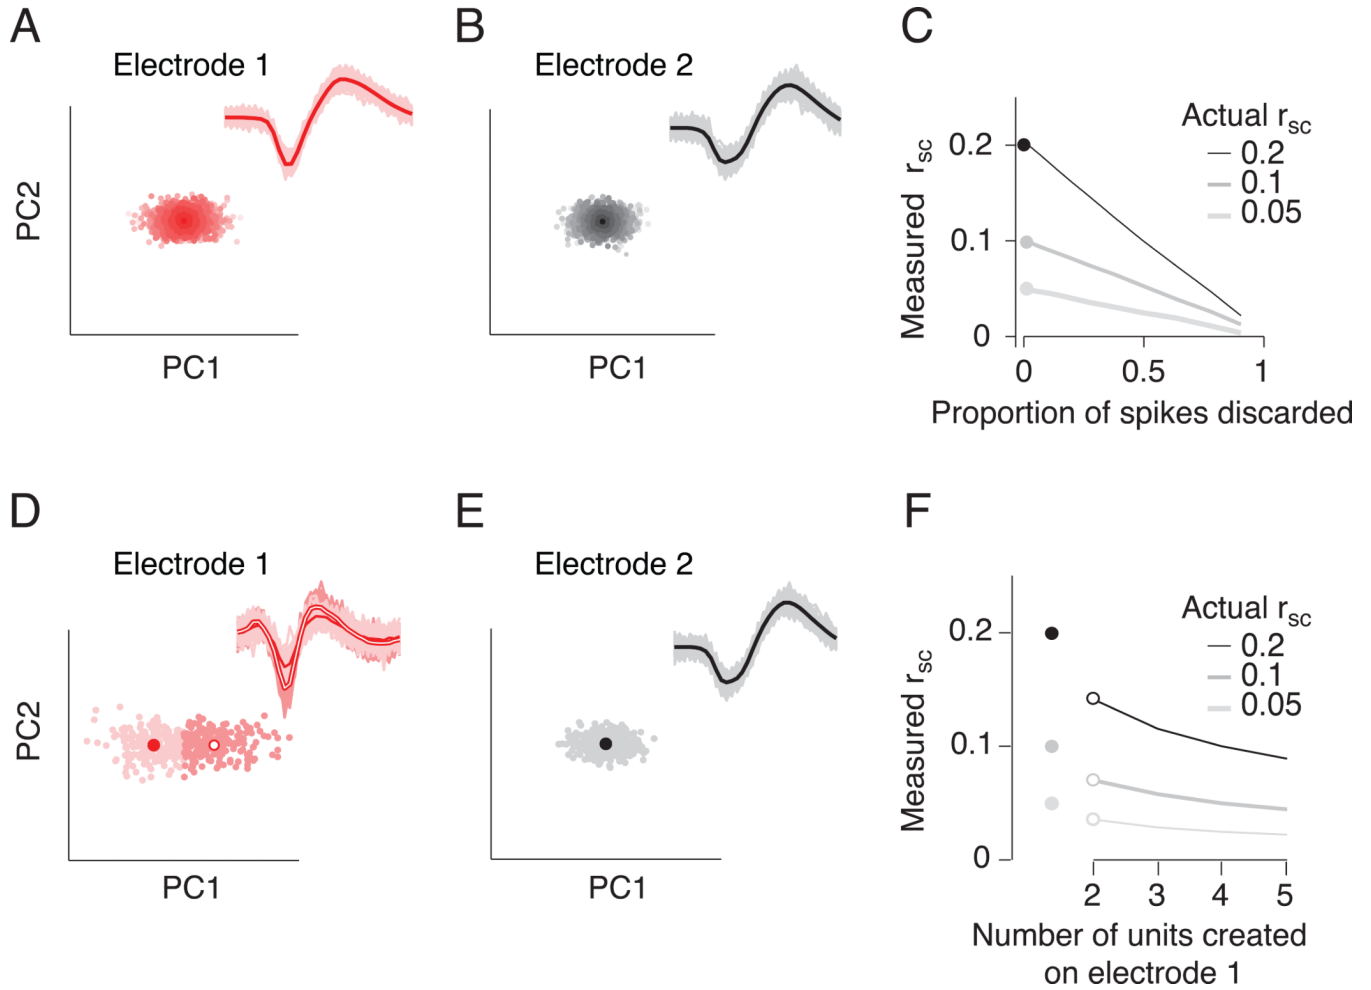


Figure 5.

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 PC space from the average waveform. **C.** Effect of increasingly stringent sorting thresholds (i.e. keeping only those spikes that are within a certain distance, in PC 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.** Same as A., 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 to 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 to 0.5 (for two units created), 0.66 (for three units created), and so on, in equation 2.




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 (e.g. 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 37% 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.

Table 1

Summary of studies measuring spike count correlations in primates. 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 listed either the average or most common value that was listed in the text or estimated from summary figures.

REFERENCE	Area	Firing rate (sp/s)	Duration (ms)	State (Task, anesthesia, etc)	rSC
Kohn and Smith (2005)*	V1	~25	2560	anesthetized	0.2
Smith and Kohn (2008)*	V1	3.4	1280	anesthetized	0.18
Reich et al (2003)	V1	1894	1894	anesthetized	0.25
Rasch et al (2011)	V1			anesthetized	0.26
Guminsky and Dragoi (2008)*	V1	~50	1860	fixation	0.25
Ecker et al (2010)*	V1	~3	500	fixation	0.01
Poort et al (2009)	V1		400	tracing	0.18
Samonds et al (2009)*	V1	30	1000	discrimination	0.1
Zandvakili and Kohn (unpublished)*	V2	10	1000	anesthetized	0.17
Smith and Sommer (2010)*	V4	5.2	1000	fixation	0.05
Cohen and Maunsell (2009)*	V4	21	200	attention/detection task	0.04
Mitchell and Reynolds (2009)*	V4	>5, ~20	800	attention/tracking task	0.05
Graf (unpublished)*	MT	~10	300	anesthetized	0.09
Huang and Lisberger (2009)*	MT	~20	500	fixation	0.1
Cohen and Newsome (2008)*	MT	28.5	500	discrimination	0.13
Zohary et al/Bair et al*	MT	~20	1000	discrimination	0.15
Erickson et al (2000)	perirhinal	~12	200-500	fixation/matching task	0.02
Averbeck and Lee (2006)	Supp motor area		66 or 200	serial reaching	0.013
Averbeck and Lee (2003)	Supp motor area	~15	200	reaching	0.02
Stark et al (2008)	Premotor areas	~5	400	grasping/imagery task	0.02
Maynard et al (1999)	M1	~20	600	reaching	0.1-0.2
Lee et al. (1998)	Motor/parietal 2/5	~5	1000	reaching	0.02-0.04
Nevet et al (2007)	Substantia nigra	58	500	cue matching	0.01-0.04
Cohen et al. (2010)	FEF	~50	few hundred	visual search	0.05-0.2
Bichot et al (2001)	FEF	~20	~200	visual search	0.09
Constantinidis/Goldman-Rakic (2002)	Prefrontal	~5	3000	delayed saccade task	0.08

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

Poort et al. (2009)⁸⁴, Samonds et al. (2009)⁸⁵, Erickson et al. (2000)⁸⁶, Averbek and Lee (2006)⁸⁷, Stark et al. (2008)⁸⁸, Maynard et al. (1999)⁸⁹, Nevet et al. (2007)⁹⁰, Cohen et al. (2010)⁹¹, Bichot et al. (2001)⁹², Khatri et al. (2009)⁹³