



Published in final edited form as:

Nat Neurosci. 2014 November ; 17(11): 1591–1597. doi:10.1038/nn.3835.

Attention can increase or decrease spike count correlations between pairs of neurons depending on their role in a task

Douglas A. Ruff and Marlene R. Cohen*

Department of Neuroscience and Center for the Neural Basis of Cognition, University of Pittsburgh, Pittsburgh, PA

Abstract

Visual attention enhances the responses of visual neurons that encode the attended location. Several recent studies showed that attention also decreases correlations between fluctuations in the responses of pairs of neurons (termed spike count correlation or r_{SC}). The previous results are consistent with two hypotheses. Attention-related changes in rate and r_{SC} might be linked (perhaps through a common mechanism), so that attention always decreases r_{SC} . Alternately, attention might either increase or decrease r_{SC} , possibly depending on the role the neurons play in the behavioral task. We recorded simultaneously from dozens of neurons in area V4 while monkeys performed a discrimination task. We found strong evidence in favor of the second hypothesis, showing that attention can flexibly increase or decrease correlations, depending on whether the neurons provide evidence for the same or opposite perceptual decisions. These results place important constraints on models of the neuronal mechanisms underlying cognitive factors.

Introduction

In recent years, correlations between the trial-to-trial fluctuations in the spiking of pairs of neurons (spike count correlations or r_{SC}) have been used to study the neuronal mechanisms underlying sensory, motor, or cognitive processes. Recent studies have shown that correlations are modulated by a wide variety of sensory, motor, and cognitive factors (for review, see ¹), most of which also alter the mean responses of cortical neurons. While a wide variety of theoretical models can typically account for mean responses, the pattern of correlation changes can place constraints on hypotheses about the underlying neural mechanisms.

Visual attention, which improves perception of an attended location or feature, is a particularly well-studied example of a cognitive process that changes rates and correlations. Attention has long been known to increase the average responses of neurons in visual cortex whose tuning matches the attended location or feature ^{2–5}. Recently, several studies have demonstrated that attention can decrease spike count correlations ^{6–10}. We wondered whether attention-related increases in rate and decreases in correlation are fixed signatures

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use:http://www.nature.com/authors/editorial_policies/license.html#terms

*Correspondence: Marlene Cohen, 4400 Fifth Ave., MI Room 115, Pittsburgh, PA 15213, Tel: 412-268-4486, cohenm@pitt.edu.

A supplementary methods checklist is available.

of a single underlying mechanism or whether attention might be associated with either increases or decreases in correlations under different task demands.

Because attention can dramatically affect perception, we reasoned that the best chance of observing attention-related increases in correlation might occur in situations in which increasing correlations would improve the ability of the population of neurons to encode information about the visual world. The relationship between correlations and information coding is complicated and an area of active investigation^{11–13}. However, we used the predictions of an influential “pooling” model of perceptual decision-making^{14,15} to create a task in which attention might reasonably be expected to increase correlations between some pairs of neurons.

In most models of decision-making, positive correlations can be either helpful or harmful. For example, the pooling model predicts that decisions between two choices are made by comparing the pooled responses of groups of neurons whose responses provide evidence in favor of each choice^{14,15}. The model predicts that positive correlations between neurons whose responses provide evidence in favor of the same choice are harmful. They reduce the benefit of averaging the responses of many neurons because shared variability cannot be averaged out. In contrast, positive correlations between neurons whose responses provide evidence for opposite choices are helpful because their shared variability can be subtracted out. Although negative correlations are not extremely common in visual cortex, their effect on population coding is thought to be the opposite of positive correlations. Negative correlations between neurons whose responses provide evidence for the same choice may be helpful because averaging their responses effectively subtracts out shared variability, whereas negative correlations between neurons whose responses provide evidence for opposite choices may be harmful because shared variability is effectively averaged when the responses of the two groups are compared.

In all studies that previously measured attention-related decreases in r_{SC} ^{6–10}, the responses of nearly all pairs of neurons for which correlations were measured contributed to the same behavioral choice. For example, in the detection task used by Cohen and Maunsell^{6,7}, monkeys were trained to report a subtle change in the orientation or spatial frequency of a flashing visual stimulus. The responses of nearly all (92%) of the neurons that they recorded from visual area V4 increased when the stimulus changed (presumably due to a release of adaptation to repeated presentations of the original stimulus). Therefore, the correlation decrease could in principle contribute to attention-related improvements in perception by improving the amount of information encoded by the population.

We hypothesized that we might observe attention-related increases in correlations in a discrimination task. Discrimination tasks cannot usually be solved by simply averaging the responses of all neurons. Instead, most hypothesized readout mechanisms compare, or take the difference between, the responses of groups of neurons providing evidence in favor of each choice. As in the detection task, correlations between neurons that provide evidence supporting the same behavioral choice may limit the benefit of averaging the responses of many neurons. If, however the total amount of neuronal variability stays the same or decreases with attention^{6,8}, positive correlations between neurons that encode opposite

choices might improve discrimination because noise that is common to the two groups will be subtracted out.

To determine whether attention-related modulations of rate and correlation are dissociable, we recorded simultaneously from a few dozen neurons in each hemisphere of visual area V4 while monkeys performed a contrast discrimination task with a spatial attention component. Consistent with previous studies, we found that attention increased average firing rates and decreased r_{SC} among pairs of neurons whose responses contribute to the same choice. However, we found that attention increased r_{SC} between pairs of neurons whose responses provide evidence in favor of opposite choices.

Our results show that directing attention to the receptive fields of pairs of visual neurons can either increase or decrease correlations. Our findings suggest that the neuronal mechanism underlying shifts in attention must affect the firing rates and correlations between visual neurons depending on the role the neurons play in the behavioral task.

Results

Contrast discrimination task and psychophysical results

We trained two monkeys (*Macaca mulatta*; both male, 7.5 and 9 kg) to perform the contrast discrimination task illustrated in Figure 1a. A trial began when the animal fixated a central spot of light. After a random period (200–400 ms, picked from a uniform distribution on each trial), two pairs of grating stimuli appeared, with one pair in each hemifield. The animals were cued in alternating blocks of trials as to which stimulus pair to base their decision on (i.e. which stimulus pair to pay attention to). After a randomly selected stimulus viewing period (333, 500, 667, or 800 ms), the gratings were replaced by two saccade targets at the locations of the gratings on the cued side. The animals' task was to make an eye movement to the target corresponding to the stimulus on the attended side that had higher contrast.

The animals were cued as to which stimulus pair to attend to in a set of separate instruction trials before each block of trials (Figure 1b). These trials were identical to the contrast discrimination trials, except that only a single grating at one of the two locations on the cued side was shown. The location of these stimuli served as a cue to the animal as to the attention condition for the upcoming trials when two pairs were shown. The single stimulus was presented at either one of the contrasts used in the contrast discrimination task or at 100% contrast, and we used these trials to characterize the neuronal responses we recorded to stimuli at different contrasts and locations. Each block of trials consisted of 40 instruction trials (5 trials at each of 2 locations and 4 contrasts) and 160 trials of the contrast discrimination task (10 per stimulus condition). During each experimental session, the animals completed at least two blocks of trials for each attention condition (at least 10 instruction trials per location–contrast and 20 contrast discrimination trials per stimulus and attention condition).

The nature of this discrimination task did not allow for catch trials to evaluate the animal's attentional state. Because of this, it was critical to balance the contrasts of all four stimuli to

prevent the animals from adopting a strategy other than basing their decisions on the relative contrasts of the two stimuli in the cued hemifield. The stimulus conditions are summarized in Figure 1c. Each pair of stimuli contained one stimulus at a fixed medium contrast (15% contrast for Monkey F and 25% for Monkey J). The other stimulus was lower on half the trials and higher on the other half. The higher and lower contrasts remained constant throughout each daily experimental session and were picked such that the contrast discrimination (medium vs. high or low contrast) was near the animal's psychophysical threshold. The stimulus pair in the uncued hemifield also contained the medium contrast and one of the same higher or lower contrasts as in the cued hemifield. On any given trial, the medium contrast occurred in exactly one stimulus in each pair and could occur at either location within a pair. The other stimulus in the uncued hemifield could either be the same (e.g. both high contrast or both low contrast) or different than the stimulus in the cued hemifield, and all trial types were randomly interleaved. There were therefore 16 unique combinations of four stimuli, each of which occurred in both attention conditions. Because the stimulus contrasts were completely balanced across hemifields and across attention conditions, attending to the stimuli in the uncued hemifield conferred no advantage to the animal.

Accordingly, behavioral evidence suggests that the animals made decisions based only on the stimuli in the cued hemifield. The animals made saccades to the locations of the stimuli in the uncued hemifield in fewer than 1.5% of trials (mean = 1.48%, median = 0.9%). This result was expected because there were no saccade targets in the uncued hemifield. We also reasoned that if the animals were not attending correctly, then the contrast of the stimuli in the uncued hemifield would have affected their decisions. Figure 1d shows that the animals' proportion correct was not significantly affected by whether the stimuli in the uncued hemifield were the same or different contrasts or higher or lower contrast as the stimuli in the cued hemifield (binomial tests, $p > 0.05$).

Neurophysiological recordings and task tuning similarity

To measure attention-related modulation of rates and correlations, we recorded from pairs of chronically implanted microelectrode arrays, one in each hemisphere of visual area V4 (48 electrodes per array, 96 per animal). The centers of the visual receptive fields from the units recorded during a representative experimental session are plotted in Figure 1e. We recorded single- and multi-unit activity from these arrays during daily experimental sessions for several weeks in each animal. Using these methods, it is nearly impossible to tell whether we recorded from the same single- or multi-unit clusters on subsequent days. To be conservative, the analyses presented in Figures 4 and 6 are based on a single recording session from each animal (picked using behavioral metrics; see Methods), but the results from the other experimental sessions were qualitatively similar (see Figure 5). During these example recording sessions, we recorded from a total of 80 single-units and multi-unit clusters from in Monkey F and 36 single-units and multi-unit clusters in Monkey J. Across 17 recording sessions, we recorded from an average of 67 single- and multi-unit clusters in Monkey F (average of 4 single-units per day), and an average of 35 single- and multi-unit clusters (average of 1 single-unit per day) in Monkey J. The recordings in Monkey J came

primarily from a single hemisphere. We based our analyses on single–units or multi–unit clusters and use the term unit to refer to either.

Our goal was to record from pairs of units whose responses provide evidence in favor of the same stimulus choice in the contrast discrimination task (to replicate the results of previous studies) as well as from pairs whose responses provide evidence for opposite stimulus choices (to determine whether attention could be associated with increases in r_{SC}). We therefore arranged the stimuli in the contrast discrimination task so that each stimulus overlapped the receptive fields of some, but not all, of the units we recorded (Figure 2a). The three example units whose contrast response curves are plotted in Figure 2b are typical of the units we recorded in that almost all units responded more to high than low contrast stimuli^{16–19}. The units in our data set typically responded most to one of the stimuli in the contralateral hemifield, and they did not respond substantially to the stimuli in the ipsilateral hemifield.

To distinguish between our two hypotheses about the relationship between attention–related modulation of rates and r_{SC} , we needed to sort unit pairs by whether their responses provided evidence in favor of the same or opposite choices in the contrast discrimination task. We therefore devised a measure of the task tuning similarity (TTS) of each pair of units recorded simultaneously in the same hemisphere. TTS quantifies the similarity in their relative responses to the stimuli at each of the two locations in the contralateral hemifield. We computed a standard d' metric for each unit comparing the distributions of its responses to 100% contrast stimuli at the two locations in the contralateral hemifield (locations 1 and 2 in the examples in Figure 2c). We arbitrarily assigned positive d' values to units that preferred location 1 over location 2 or location 3 over location 4 and negative d' values to the units with opposite preference. To compute task tuning similarity, we simply multiplied the d' values for each pair of units. Therefore, pairs with positive TTS responded preferred the same stimulus location and pairs with negative TTS preferred opposite locations.

We could in principle have defined task tuning similarity in many ways, and we have no way of knowing the role that each of the units we recorded actually played in the animal's decision. We found that r_{SC} varies systematically along our TTS axis (see analyses below), which shows that this measure captures some aspects of the relationship between the tuning properties of the neurons we recorded and their attention–related modulation.

Attention can either increase or decrease r_{SC}

The predictions of our two hypotheses can be easily distinguished as a function of TTS. Studies throughout visual cortex have shown that pairs of units with similar receptive fields or tuning properties tend to have higher r_{SC} than those with dissimilar tuning^{6,14,20–29}. Because TTS measures the extent to which two units prefer the same or opposite stimulus locations, pairs with positive TTS will tend to have more similar receptive fields than pairs with negative TTS. Both hypotheses therefore predict that r_{SC} will increase with TTS while the animal is ignoring the stimuli in the contralateral hemifield (i.e., the red lines have positive slopes in Figures 3a and 3b).

Both hypotheses predict that, consistent with previous studies^{6–10}, attention will decrease r_{SC} for pairs whose responses provide evidence for the same choice (positive TTS, "same group"). The first hypothesis predicts that attention will decrease correlations regardless of TTS, so r_{SC} should always be lower when the animal is discriminating the stimuli in the contralateral hemifield (green line in Figure 3a). The second hypothesis (Figure 3b) predicts that attention will have a qualitatively different effect on pairs of units depending on the sign of their TTS, decreasing r_{SC} when TTS is positive ("same group") and increasing r_{SC} when TTS is negative ("different group").

The results strongly support the second hypothesis and show that attention can either increase or decrease r_{SC} . Consistent with previous results, attention reduced correlations among pairs of units with positive TTS. In both animals, however, attention increased r_{SC} among pairs with negative TTS (Figures 4a and 4b, 3124 same-hemisphere pairs in Monkey F and 854 same-hemisphere pairs in Monkey J for these example recording sessions). In this example data set in Monkey J, the average correlation coefficient for neurons with $TTS < -1$ was less than 0. Although there was considerable variability in r_{SC} across neuron pairs and recording sessions, an average r_{SC} less than 0 was not extremely common across our data set (see Figure 5). Consistent with our average results, studies in visual cortex typically report small, but positive average correlations¹ although studies in frontal cortex have noted consistently negative correlations, particularly in pairs of neurons whose receptive fields do not overlap^{30,31}. Simple circuit models can produce both positive and negative correlations simply from variability in the activity and strength of inhibitory and excitatory inputs³², and the magnitude of correlations depends on a wide variety of experimental and theoretical factors¹. We therefore focus here on attention-related changes in spike count correlations, rather than their absolute magnitude.

The different effects of attention on correlations between neurons with positive and negative TTS was consistent across recording sessions. Across 17 recording sessions (6 from monkey F and 11 from monkey J), attention was associated with increased firing rates both for units that contributed to pairs with high positive or negative TTS (Figure 5 left and right, $TTS < -1$ and $TTS > 1$, respectively). These increases were on the small end of the range reported in previous studies, perhaps because having four stimuli on the screen meant that both the receptive field center and surround were stimulated for most neurons. However, even though attention was associated with similar firing rate changes for neurons that contributed to pairs with $TTS < -1$ or $TTS > 1$, attention had opposite effects on r_{SC} for these pairs, increasing correlations among pairs with $TTS < -1$ (Figure 5 left) and reducing correlations among pairs of units with $TTS > 1$ (Figure 5, right).

The impact of attention on r_{SC} can be described as decreasing the slope of the line relating r_{SC} to TTS. To quantify this effect, we fit a line to scatter plots of r_{SC} as a function of TTS in each attention condition and compared the slopes in the attended and unattended conditions. In our example recording sessions, attention was associated with a decrease in slope from 0.038 to 0.023 (Monkey F) and 0.039 to 0.005 (Monkey J). We performed a bootstrap test by randomly assigning the two correlation coefficients for each pair (one per attention condition) to new bootstrapped conditions. We then fit lines to each bootstrapped condition and calculated a distribution of the differences in slope between the two

bootstrapped conditions. The actual attention-related decreases in slope were statistically significant in 15 of our 17 data sets (bootstrap test, $p < 0.05$).

Together, our results show that attention is not obligatorily associated with decreases in correlations.

Possible artifacts

Our goal was to determine the effect of shifting attention on correlations between the responses of pairs of V4 units with different properties. In this section, we examine a number of factors that could in principle have artifactually affected our correlation results.

The first category of potential artifacts concerns the properties of the units that form the pairs with very positive and negative TTS. Although most units contribute to pairs with both positive and negative TTS, asymmetries in the number and properties of units that respond to each stimulus mean that the distributions of units that contribute to the two extremes of the x-axes in Figure 4 or separate bar graphs in Figure 5 are not identical. To be sure that our results are not confounded by any differences, we subsampled our data set to control for several factors that might affect attention-related changes in r_{SC} .

Measured correlations are known to covary with firing rate^{1,33} and also tend to be higher in multi-units than single-units¹ which comprise the majority of our data set. These factors, in addition to a longer time window over which we measured responses¹, almost certainly contribute to the relatively high spike count correlations we observed. A previous study found that attention-related modulation in rate and in r_{SC} did not differ substantially between single- and multi-units⁶, so it seems likely that we would have observed similar attention-related changes in r_{SC} if we recorded solely from single-units.

However, we performed an additional control analysis to be sure that the attention-related changes in r_{SC} were not caused by changes in rate or by asymmetries in firing rate changes between the units that contribute to pairs with positive or negative TTS. Using the distribution-matching procedure described in Methods, we calculated r_{SC} among subsets of unit pairs recorded during the two example recording sessions depicted in Figure 4 for which the four firing rate distributions for pairs of units with $TTS < -1$ or $TTS > 1$ in the two attention conditions were matched. This manipulation did not affect the qualitative changes in r_{SC} for pairs with positive or negative TTS (Figure 6, t-tests, $p < 0.05$), suggesting that differences in firing rate cannot explain our results. For comparison, the raw data from Figure 4 have been plotted and averaged in Figure 6, and the average firing rates for the pairs with $TTS < -1$ or $TTS > 1$ have been plotted in the lower panels of Figure 6.

As predicted, pairs with positive or negative TTS differed in their baseline level of r_{SC} . To control for a possible relationship between mean r_{SC} and attention-related modulation of r_{SC} (e.g. a floor or ceiling effect), we matched the distributions of r_{SC} in the attended condition for pairs with $TTS < -1$ or $TTS > 1$ for the same example data sets. In the unattended condition, r_{SC} remained lower than in the attended condition for this subset of pairs with $TTS < -1$ and higher than in the attended condition for this subset of pairs with $TTS > 1$ (Figure 6, t-tests, $p < 0.05$). Differences in mean r_{SC} therefore cannot explain our results.

Our measure of TTS depended heavily on the overlap between the spatial receptive fields of the two neurons in a pair. The dependence of the attention-related correlation changes on TTS could therefore be a hardwired signature of how cognitive factors affect neurons with overlapping or nonoverlapping receptive fields. To address this issue, we controlled for the distributions of distances between the electrodes that recorded pairs of neurons with $TTS > 1$ or $TTS < -1$. The mean electrode distance for pairs with $TTS > 1$ was significantly smaller than for pairs with $TTS < -1$ for the example recording session in Monkey F (mean distance 1.19 mm for $TTS > 1$ vs. 2.45 for $TTS < -1$, t-test, $p < 10^{-5}$) but not for Monkey J (mean distance 2.19 mm for $TTS > 1$ vs. 2.13 for $TTS < -1$, t-test, $p = 0.9$). This difference might be attributed to the more foveal receptive fields in Monkey J than Monkey F.

In both monkeys, however, the qualitatively different attention-related changes in r_{SC} remained when we matched distributions of electrode distance for pairs with $TTS > 1$ and $TTS < -1$ (right bars of Figure 6, t-tests, $p < 0.05$). Therefore, differences in mean electrode distance cannot explain the r_{SC} differences we observed. This result rules out a possible mechanism by which attention has qualitatively different effects on pairs of neurons based solely on their cortical distance.

Another category of possible artifacts do not lend themselves to data analysis controls, but aspects of our experimental design make it extremely unlikely that these qualitatively affected the pattern of attention-related changes in r_{SC} that we observed. For example, fluctuations in the animals' global internal states like arousal, alertness, or motivation could affect measurements of r_{SC} . However, these factors can be expected to affect the responses of all units and would therefore affect r_{SC} independent of TTS. Furthermore, because we recorded from both hemispheres simultaneously (with about equal numbers of units in each hemisphere in Monkey F), biases in the monkey's cognitive state (e.g. greater arousal when attending to the left) would affect both attention conditions equally, since attending to the left is the attended condition for units in the right hemisphere and the unattended condition for units in the left hemisphere.

Factors like fixational eye movements or shifts in spatial attention between the two stimuli within a hemifield could affect r_{SC} in a way that depends on TTS, but these would lead to the opposite pattern of results than the one we observed. Making small eye movements or shifting attention between the two stimuli in one hemifield in the process of discriminating between them (i.e. in the attended condition) would cause an increase in r_{SC} among units with positive TTS, since these factors would co-modulate their responses. Because neurons with negative TTS tend to have different receptive fields, these same eye movements or attentional shifts would cause a decrease in r_{SC} among pairs with negative TTS because these factors would anti-correlate their responses. For example, a small eye movement would be likely to shift a stimulus either into or out of the joint receptive field of neurons with similar receptive fields (positive TTS), but might shift the stimulus into the receptive field of one neuron but out of the receptive field of the other neuron in a pair with different receptive fields. Therefore, these factors would cause attention-related increases in r_{SC} for pairs with positive TTS and decreases for pair with negative TTS, which is the opposite of what we observed. In summary, we cannot think of an experimental artifact that could produce the pattern of attention-related changes in r_{SC} that we observed.

Discussion

Our results show that attention can either increase or decrease correlations, depending on the role the neurons play in the task. We were surprised by this result because a large body of previous work suggested that attention would always decrease correlations. The previous studies that measured attention-related modulation of r_{SC} found that attention decreases correlations^{6–10}. Other processes that increase the responses of neurons in primate visual cortex have also been shown to decrease r_{SC} , including increasing the contrast of a visual stimulus²⁰, training on a perceptual task³⁴, and the absence of adaptation²². It should be noted, however, that in two non-primate systems in which firing rates did not change, novelty or salience could either increase or decrease r_{SC} ^{35,36}.

In particular, we originally hypothesized that attention-related modulation of rates and correlations are inextricably linked because our previous study showed that in a detection task, two forms of attention (spatial attention and a form of feature attention) have the same quantitative relationship between increases in firing rate and decreases in r_{SC} ⁷. When two neurons showed a large increase in rate, they tended to show a predictable and large decrease in r_{SC} (a correlation decrease of approximately 0.05 for every 10 spikes/s of rate increase). Neurons that showed no attention-related rate change tended not to show a correlation change. This quantitative similarity between two types of attention was consistent with the hypothesis that a single neuronal mechanism underlies rate and correlation changes caused by both types of attention.

Although it is indisputable that there is a relationship between firing rate and correlation changes^{1,33}, our results show that groups of neurons with the same attention-related firing rate changes can have qualitatively different changes in correlation. Our control analyses showed that the qualitatively different attention-related changes in r_{SC} for pairs with positive or negative TTS remained even when we matched distributions of firing rates across attention and TTS conditions. Therefore, if a common mechanism is responsible for changes in both firing rates and correlations, it must affect rates and correlations in a separable way. It would be interesting to determine whether the magnitude of correlation changes of either sign are related to behavioral measures of attention. It would be difficult to design a task that involves neurons with both positive and negative TTS and also contains catch trials to allow a behavioral assessment of attention, but this would be an extremely interesting challenge for future work.

Our results complement those of a previous study of task-related changes in r_{SC}

A previous study measured how correlations between pairs of neurons in area MT depended on the role the neurons played in a direction discrimination task²³. This study measured correlations in two task conditions, which differed in both whether the neurons contributed to the same or opposite decisions (an analog of positive or negative TTS in our study) and whether the feature the neurons encoded was attended or unattended. The results of this previous study complement our current study by showing that r_{SC} depends on a combination of TTS and feature attention. In contrast, our study dissociated the effects of attention on correlations between pairs of neurons with a fixed TTS. Together, the two studies show that r_{SC} depends on the role that neurons play in the task as well as on both feature and spatial

attention. In addition, our results show that task-related changes in r_{SC} are dissociable from changes in firing rate, the baseline correlation between two neurons, and the physical distance between the two neurons in the brain.

Relationship between correlations and information coding

A growing body of correlation data from different tasks and experimental systems shows that correlations depend on much more than hard-wired feedforward common inputs. A tantalizing possibility is that attention and other sensory, motor, and cognitive processes can modulate rates and correlations in the way that is best suited for information coding. It is currently impossible to prove whether this is generally true for both experimental and theoretical reasons. First, all available data come from animals that have been overtrained on specific tasks, leaving open the possibility that attention produces beneficial changes in correlations only in very specific instances. Second, the impact of correlations on information coding is a matter of current debate in the computational community^{12,13,15}.

However, at least in one simple framework, our results seem consistent with the idea that attention modulates correlations to optimize information coding. Attention, like many other sensory, motor, or cognitive processes, typically multiplicatively scales the responses of visual neurons³⁷. Such increases in response gain are always good for information coding because they improve the signal to noise ratio of single neurons. Changes in r_{SC} have the potential to have an even bigger impact on information coding than changes in single neurons because correlations can, according to some models, either limit or improve the benefit of pooling information over many neurons^{6,8,13–15}. In our study, attention affected r_{SC} in a manner consistent with improving information coding by decreasing correlations for pairs with positive TTS and increasing correlations among pairs with negative TTS.

How the neuronal mechanisms underlying processes like attention could selectively modulate correlations for the purpose of improving information coding remains an open question. Our results show that any potential mechanism must at the very least be flexible enough to dissociate changes in correlations from changes in response rate, baseline correlation, or cortical distance.

Methods

The subjects in our experiment were two adult male rhesus monkeys (*Macaca mulatta*, 7.5 and 9 kg). All animal procedures were approved by the Institutional Animal Care and Use Committees of the University of Pittsburgh and Carnegie Mellon University. Before training, we implanted each animal with a titanium head post. After the animal learned the task (5–7 months), we implanted a pair of 6×8 microelectrode arrays (Blackrock Microsystems), one area V4 in each cerebral hemisphere. The distance between adjacent electrodes was 400 μm , and each electrode was 1 mm long. We identified area V4 using stereotactic coordinates and by visually inspecting the sulci. We placed the arrays between the lunate and the superior temporal sulci. The two arrays were connected to a single percutaneous connector that allowed electrophysiological recordings.

We recorded neuronal activity from these arrays during daily experimental sessions for several weeks in each animal. Using our recording methods, it is nearly impossible to tell whether we recorded from the same single- or multi-unit clusters on subsequent days. To be conservative, the analyses presented here are based on a single recording session from each animal. These example days were picked because the animal performed a large number of trials with good psychophysical performance and because recording quality was good.

We confirmed that the results from these example recording sessions were typical of our data set by analyzing data from all recording sessions in which the animal completed at least 250 contrast discrimination trials at each attended location (median 1079 correct trials across both attended locations), achieved at least 60% correct performance during contrast discrimination trials, made choices toward one location in a hemifield no more than 2.5 times as often as the other location, and had good recording quality (we successfully recorded from most electrodes, the recordings were largely free from electrical noise and the stimuli were appropriately placed over the units' receptive fields). Seventeen recording sessions fulfilled all of these criteria (6 from monkey F and 11 from monkey J). In these 17 sessions, the mean percent correct was 69.5% (range 62%–83%). The monkeys performed above chance during every session (binomial tests, $p < 10^{-3}$). In the example sessions we analyzed more in depth, performance was 82% correct and 73% correct for Monkeys F and J, respectively.

We recorded a total of 34 single-units and 805 multi-unit clusters across these sessions. Because it is impossible to know the extent to which the same units were recorded on subsequent days, the error bars in Figure 6 should be considered an upper bound on statistical significance. When we treated the observations as independent, t-tests showed that the attention-related changes in firing rate and correlation were highly significant, but assumption of independence is almost certainly incorrect. The analysis in Figure 6 shows that our example recording sessions were typical of our data set, but these data analysis concerns led us to conservatively base all conclusions on statistical analyses done within the example recording sessions for each animal.

All spike sorting was done manually following the experiment using Plexon's Offline Sorter. We sorted single-units as well as multi-unit clusters whose waveforms looked like action potentials and whose distributions of interspike intervals looked plausibly neural. We included single-units or multi-unit clusters for analysis if their response from 50 to 100 ms after stimulus onset (averaged over all stimuli) was significantly different than its baseline firing rate sampled 50 ms before stimulus onset (t-test, $p < 0.05$).

We presented visual stimuli on a calibrated CRT monitor (calibrated to linearize intensity; 1024×768 pixels; 120 Hz refresh rate) placed 54 cm from the animal. We monitored eye position using an infrared eye tracker (Eyelink 1000 (SR Research, Ontario, Canada)). We used custom software (written in Matlab using the Psychophysics Toolbox^{38,39}) to present stimuli and monitor behavior. We recorded eye position and pupil diameter (1000 samples/s), neuronal responses (30,000 samples/s) and the signal from a photodiode to align neuronal responses to stimulus presentation times (10,000 samples/s) using hardware from Ripple (Salt Lake City, UT).

Data analysis

To allow for the latency of V4 responses, our analyses are based on spike count responses calculated from 60 ms to 393 ms after stimulus onset. We quantified spike count correlations (r_{SC}) as the Pearson's correlation coefficient between spike count responses to repeated presentations of the same stimulus. This measure is extremely sensitive to outliers, so we did not analyze trials for which the response of either unit was more than three standard deviations away from its mean (following the convention in ²⁰). For each pair of units recorded simultaneously from the same hemisphere but not from the same electrode, we computed r_{SC} separately for each stimulus condition and averaged the results. Taking the z-scored responses for each condition and computing a single value of r_{SC} for each pair (as in ²⁴) gave qualitatively similar results.

The controls in Figure 6 are based on matched distributions of units or pairs. The goal of these controls was determine whether the attention-related changes in r_{SC} that we observed could be attributed to differences in either mean rate or baseline r_{SC} across attention conditions or between pairs with positive or negative TTS. We therefore subsampled our data to create subdistributions of pairs such that the distributions of either firing rate or r_{SC} in the attended condition were identical. For firing rate, we matched four total distributions (attended and unattended for pairs with $TTS < -1$ or $TTS > 1$). For baseline r_{SC} , we matched two distributions ($TTS < -1$ or $TTS > 1$ in the attended condition).

To create matched distributions, we first binned the data to create histograms of the distributions of geometric mean firing rate in a pair or of r_{SC} in the attended condition. We then picked, without replacement, from each bin of each distribution to create subdistributions in which each subdistribution has an identical number of points in a given bin. For example, for each bin, we would look to see which of the four original firing rate distributions has the fewest data points in that bin, and we would pick a random subset of the data from the other distributions such that each of the four new distributions has the same number of data points in that bin. We repeated this resampling procedure 10,000 times, and the numbers in Figure 6 represent the average of these resampled distributions. The error bars represent the standard error for a representative resampled distribution.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Regina Chang for assistance with animal training and recordings, Karen McCracken for technical assistance, and John Maunsell and Brent Doiron for comments on an earlier version of this manuscript. The authors are supported by NIH grants 4R00EY020844-03 and R01 EY022930 (MRC), a training grant slot on NIH 5T32NS7391-14 (DAR), a Whitehall Foundation Grant (MRC), a Klingenstein Fellowship (MRC) and a Sloan Research Fellowship (MRC).

References

1. Cohen MR, Kohn A. Measuring and interpreting neuronal correlations. *Nat. Neurosci.* 2011; 14:811–819. [PubMed: 21709677]

2. Maunsell JHR, Cook EP. The role of attention in visual processing. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* 2002; 357:1063–1072. [PubMed: 12217174]
3. Maunsell JHR, Treue S. Feature-based attention in visual cortex. *Trends Neurosci.* 2006; 29:317–322. [PubMed: 16697058]
4. Reynolds JH, Chelazzi L. Attentional modulation of visual processing. *Annu. Rev. Neurosci.* 2004; 27:611–647. [PubMed: 15217345]
5. Yantis S, Serences JT. Cortical mechanisms of space-based and object-based attentional control. *Curr. Opin. Neurobiol.* 2003; 13:187–193. [PubMed: 12744972]
6. Cohen MR, Maunsell JHR. Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* 2009; 12:1594–1600. [PubMed: 19915566]
7. Cohen MR, Maunsell JHR. Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron.* 2011; 70:1192–1204. [PubMed: 21689604]
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. Herrero JL, Gieselmann M, Sanayei M, Thiele A. Attention-induced variance and noise correlation reduction in macaque V1 is mediated by NMDA receptors. *Neuron.* 2013; 78:729–739. [PubMed: 23719166]
10. Zénon A, Krauzlis R. Attention deficits without cortical neuronal deficits. *Nature.* 2012; 489:434–437. [PubMed: 22972195]
11. Averbeck BB, Latham PE, Pouget A. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* 2006; 7:358–366. [PubMed: 16760916]
12. Ecker AS, Berens P, Tolias AS, Bethge M. The effect of noise correlations in populations of diversely tuned neurons. *J. Neurosci.* 2011; 31:14272–14283. [PubMed: 21976512]
13. Abbott LF, Dayan P. The effect of correlated variability on the accuracy of a population code. *Neural Comput.* 1999; 11:91–101. [PubMed: 9950724]
14. Zohary E, Shadlen M, Newsome W. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature.* 1994; 370:140–143. [PubMed: 8022482]
15. Shadlen MN, Britten KH, Newsome WT, Movshon JA. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* 1996; 16:1486–1510. [PubMed: 8778300]
16. Dean A. The Variability of Discharge of Simple Cells in the Cat Striate Cortex. *Exp. Brain Res.* 1981; 44:437–440. [PubMed: 7308358]
17. Albrecht DG, Hamilton DB. Striate Cortex of Monkey and Cat: Function Contrast Response. *J. Neurophysiol.* 1982; 48:217–237. [PubMed: 7119846]
18. Sclar G, Freeman RD. Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Exp. Brain Res.* 1982; 46:457–461. [PubMed: 7095050]
19. Sclar G, Maunsell J, Lennie P. Coding of image contrast in central visual pathways of the macaque monkey. *Vision Res.* 1990; 30:1–10. [PubMed: 2321355]
20. Kohn A, Smith MA. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* 2005; 25:3661–3673. [PubMed: 15814797]
21. Smith MA, Kohn A. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* 2008; 28:12591–1603. [PubMed: 19036953]
22. Gutnisky, Da; Dragoi, V. Adaptive coding of visual information in neural populations. *Nature.* 2008; 452:220–224. [PubMed: 18337822]
23. Cohen MR, Newsome WT. Context-dependent changes in functional circuitry in visual area MT. *Neuron.* 2008; 60:162–173. [PubMed: 18940596]
24. Ecker AS, et al. Decorrelated Neuronal Firing in Cortical Microcircuits. *Science (80-.).* 2010; 327:584–587.
25. 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]

26. Jermakowicz W, Chen X, Khaytin I, Bonds A, Casagrande V. Relationship between spontaneous and evoked spike-time correlations in primate visual cortex. *J. Neurophysiol.* 2009; 101:2279–2289. [PubMed: 19211656]
27. Smith MA, Jia X, Zandvakili A, Kohn A. Laminar dependence of neuronal correlations in visual cortex. *J. Neurophysiol.* 2013; 109:940–947. [PubMed: 23197461]
28. Hansen BJ, Chelaru MI, Dragoi V. Correlated variability in laminar cortical circuits. *Neuron.* 2012; 76:590–602. [PubMed: 23141070]
29. Smith MA, Sommer MA. Spatial and temporal scales of neuronal correlation in visual area V4. *J. Neurosci.* 2013; 33:5422–5432. [PubMed: 23516307]
30. Cohen JY, et al. Cooperation and competition among frontal eye field neurons during visual target selection. *J. Neurosci.* 2010; 30:3227–3238. [PubMed: 20203182]
31. Leavitt ML, Pieper F, Sachs A, Joobar R, Martinez-Trujillo JC. Structure of spike count correlations reveals functional interactions between neurons in dorsolateral prefrontal cortex area 8a of behaving primates. *PLoS One.* 2013; 8:e61503. [PubMed: 23630595]
32. Renart A, et al. The asynchronous state in cortical circuits. *Science (80-.).* 2010; 327:587–590.
33. De la Rocha J, Doiron B, Shea-Brown E, Josi K, Reyes A. Correlation between neural spike trains increases with firing rate. *Nature.* 2007; 448:802–806. [PubMed: 17700699]
34. Gu Y, et al. Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron.* 2011; 71:750–761. [PubMed: 21867889]
35. Miura K, Mainen Z, Uchida N. Odor representations in olfactory cortex: distributed rate coding and decorrelated population activity. *Neuron.* 2012; 74:1087–1098. [PubMed: 22726838]
36. Jeanne JM, Sharpee TO, Gentner TQ. Associative Learning Enhances Population Coding by Inverting Interneuronal Correlation Patterns. *Neuron.* 2013; 78:352–363. [PubMed: 23622067]
37. McAdams CJ, Maunsell JHR. Effects of Attention on Orientation-Tuning Functions of Single Neurons in Macaque Cortical Area V4. *J. Neurosci.* 1999; 19:431–441. [PubMed: 9870971]
38. Brainard DH. The Psychophysics Toolbox. *Spat. Vis.* 1997; 10:433–436. [PubMed: 9176952]
39. Pelli D. The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spat. Vis.* 1997; 10:437–442. [PubMed: 9176953]

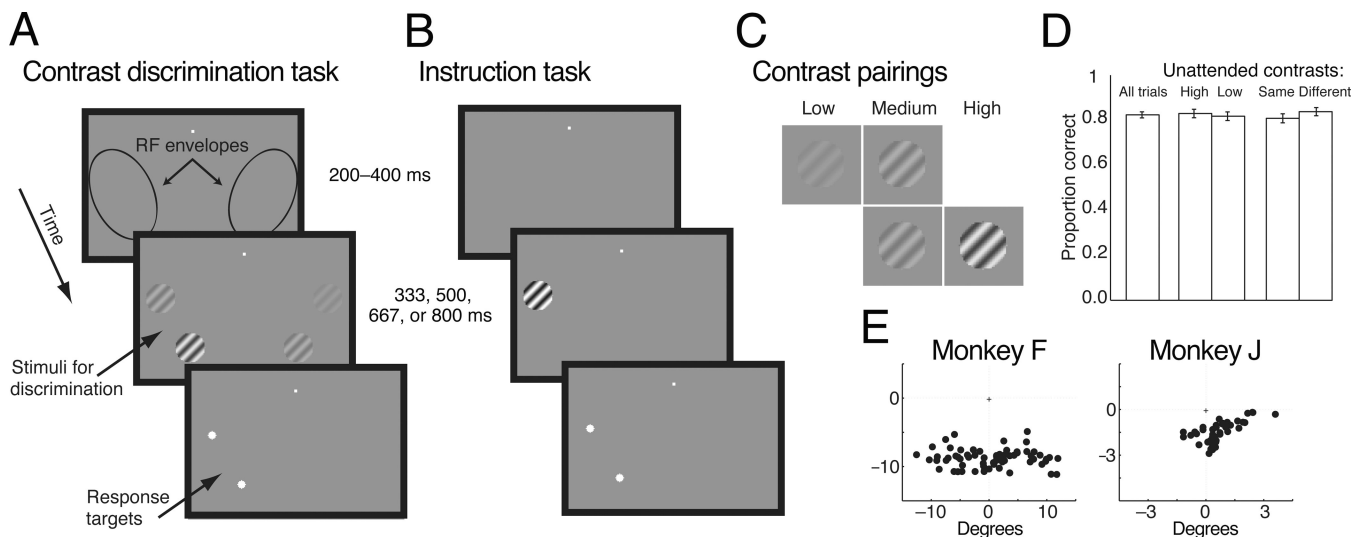


Figure 1. Task and stimuli

(a) Contrast discrimination task. The animals were required to maintain fixation in a $1\text{--}1.5^\circ$ window for the duration of the trial. After the animal fixated a blank screen for 200–400 ms, two pairs of grating stimuli were placed on each side of a fixation spot within the envelope of the receptive fields of the units recorded in each hemisphere. The animal's task was to determine which of the pair of stimuli in the previously cued hemifield had higher contrast. After the stimulus viewing time (333, 500, 667, or 800 ms), two targets appeared at the locations of the stimuli in the cued hemifield. When the targets were presented, the animal was free to move its eyes and was rewarded for saccades directed to the corresponding to the higher contrast stimuli in the pair. (b) Instruction trials. At the start of each attention block, the animal performed instruction trials in which only a single stimulus was presented at one of the stimulus locations in a single hemifield, were presented. The location of the stimulus served as a cue to the animal as to the hemifield to attend to in the following block of trials and provided data we used to assess the unit's contrast response function at each of the four locations. The timing of instruction trials was identical to the timing in the contrast discrimination trials, and the animal was rewarded for saccades directed to the target at the location of the single stimulus. (c) Discrimination task contrast pairings. During the discrimination task, each pair of stimuli always contained a medium contrast stimulus. This stimulus could appear at either location within the pair. The second stimulus in each pair was either a higher or lower contrast stimulus. There were 16 unique combinations of the four stimuli, and these were randomly interleaved in each block of trials. (d) Performance did not depend on the stimuli in the opposite hemifield. This figure depicts average proportion correct during the example recording sessions from Monkeys F and J whose physiological results are described in Figures 5 and 7. Performance did not significantly depend on whether the contrast of the unattended stimulus that was not medium contrast was high or low or whether this stimulus was the same or a different contrast from the contrasts of the attended stimuli (binomial test, $p < 0.05$). Error bars are standard errors of the mean (s.e.m). (e) Center of the visual receptive fields for the multi-unit signals from example recording sessions for each animal.

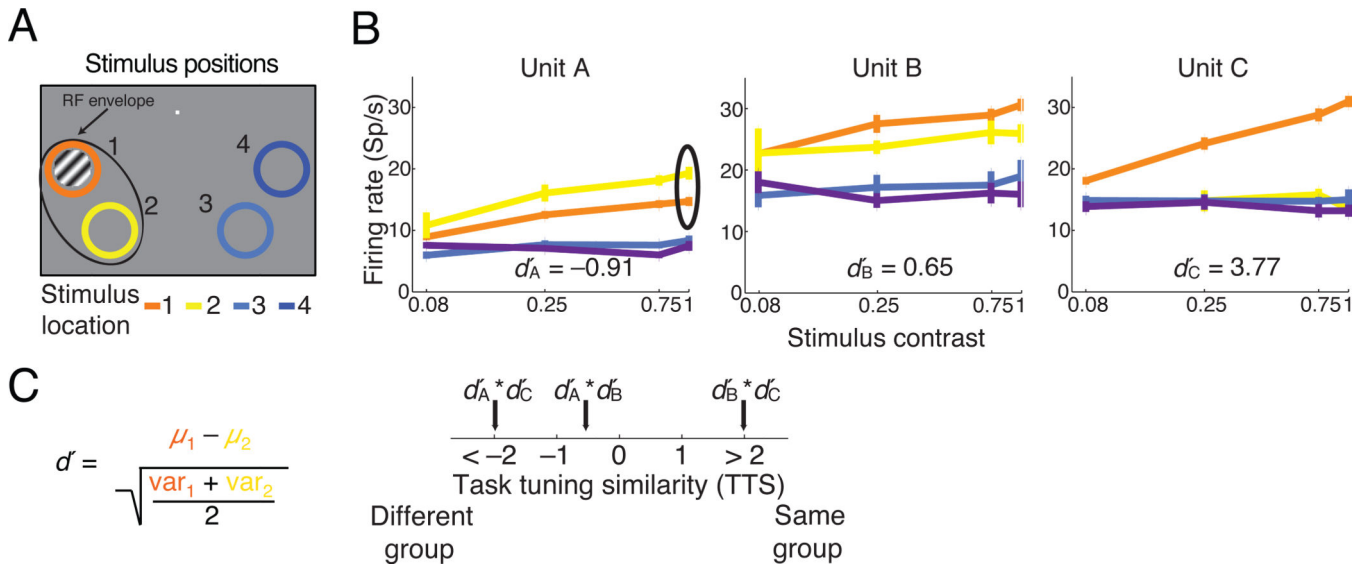


Figure 2. Task tuning similarity calculation

(a) Cartoon of an example instruction trial. The colored rings represent the locations of the four stimuli during the contrast discrimination task. The rings were not visible to the monkey and serve only as a key for panel b. (b) Contrast response functions generated from instruction trials for three example units (error bars represent s.e.m.). We quantified the unit's preference for each of the two locations in the contralateral hemifield by computing a signed d' from its responses to 100% contrast stimuli. (c) The Task Tuning Similarity (TTS) axis represents the product of the d' s calculated for each of pair of units recorded simultaneously within the same hemisphere. Positive TTS indicates that both units in a pair prefer the same stimulus location, and negative products signify opposite location preferences.

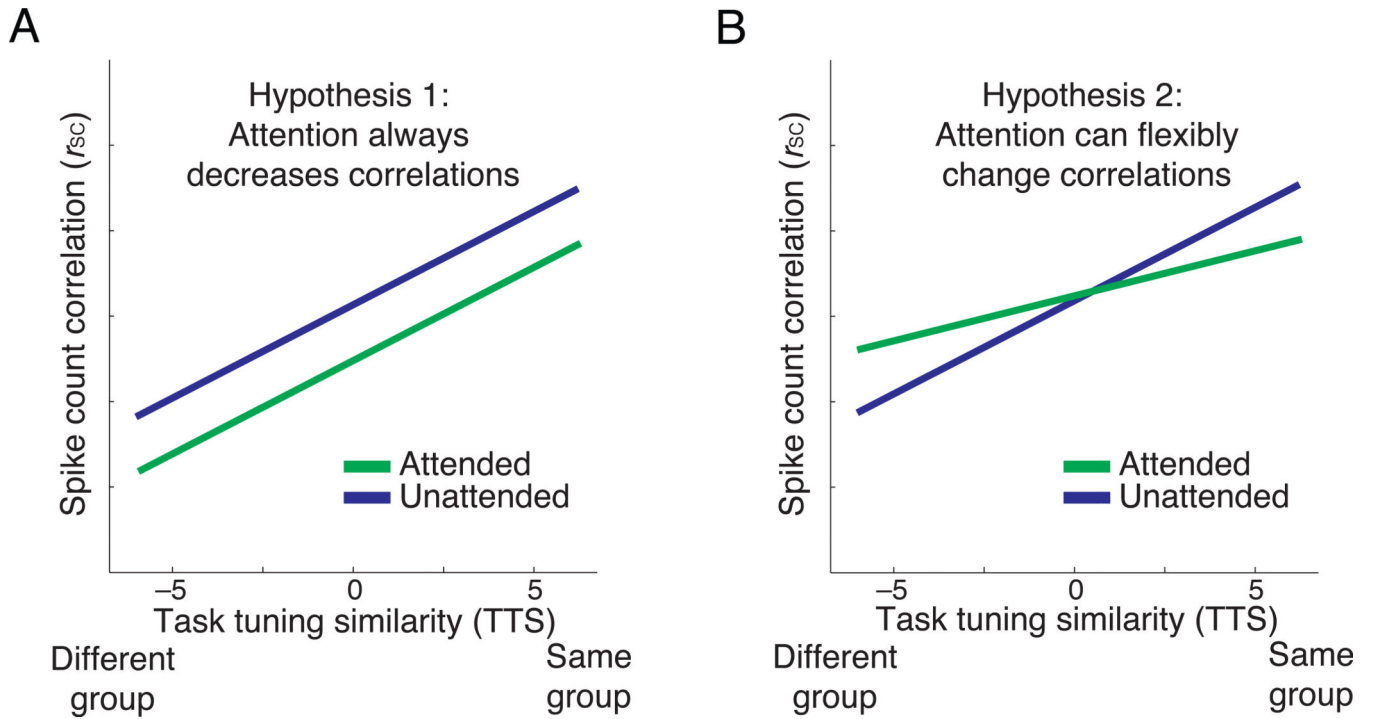


Figure 3. Hypotheses for how spike count correlations depend on TTS

(a) Hypothesis 1: attention decreases correlations regardless of the relationship between a pair of unit's tuning similarity. (b) Hypothesis 2: attention can either increase or decrease correlations. Under one model, if attention changes correlations to maximize information coding, correlations would be expected to decrease between pairs with positive TTS but increase between pairs with negative TTS.

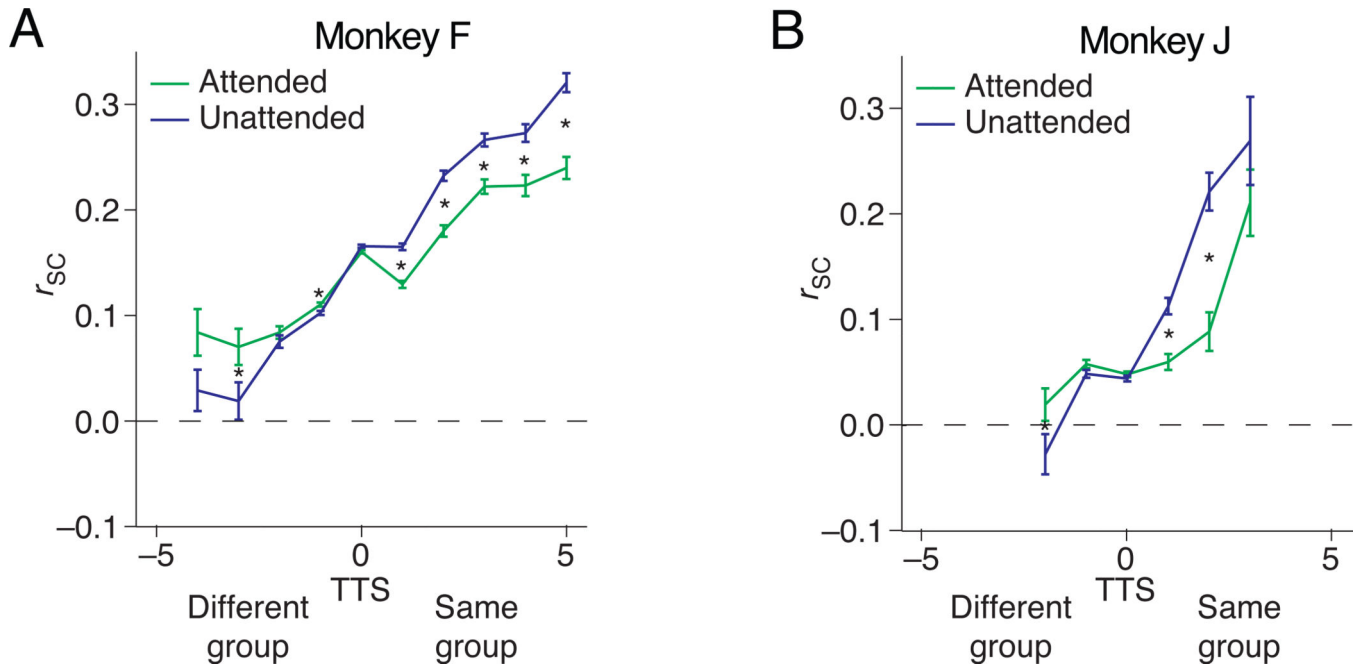


Figure 4. Attention can either increase or decrease spike count correlations

(a) and (b) depict pairwise correlation values from an example recording session from Monkeys F and J, respectively. Consistent with the second hypothesis in Figure 3, attention decreased correlations amongst pairs of units with positive TTS and increased correlations between units with negative TTS. Error bars are s.e.m. The bins are non-overlapping groupings of pairs with different TTS (bin width = 1). Asterisks indicate bins for which the r_{SC} was significantly different in the two attention condition (paired t-test on the z-transformed correlation coefficients; $p < 0.05$ after Bonferroni correction). The number of pairs contributing to each bin is, from left, 14, 16, 64, 756, 1182, 206, 82, 44, 20, and 22 for Monkey F and 38, 306, 580, 78, 22, and 13 for Monkey J.

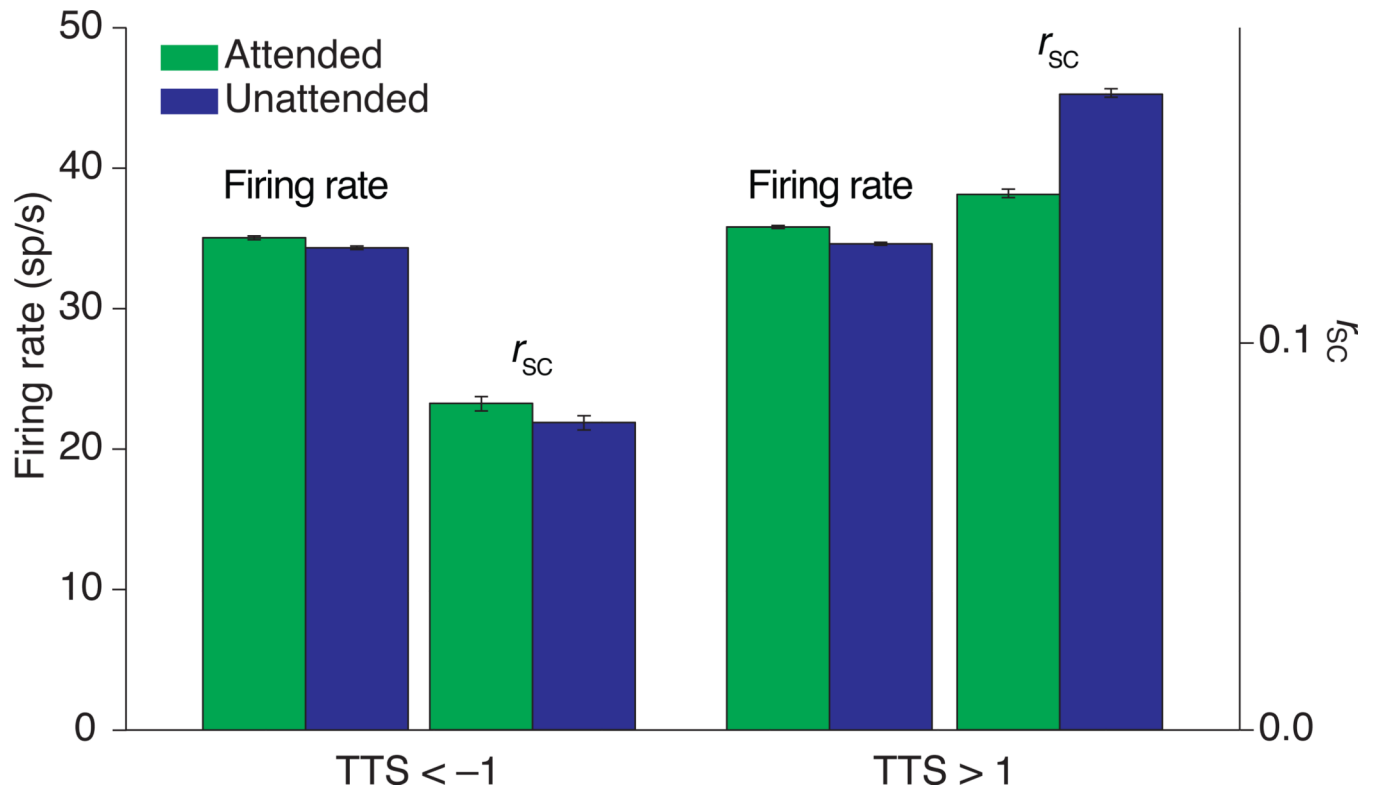


Figure 5. Summary of results across 17 recording sessions
 Average firing rate and spike count correlations when attention was directed to the contralateral (attended) and ipsilateral (unattended) hemisphere for units contributing to pairs with TTS < -1 (left) and TTS > 1 (right). Error bars are s.e.m.

