

Attention Increases Spike Count Correlations between Visual Cortical Areas

 Douglas A. Ruff and  Marlene R. Cohen

Department of Neuroscience and Center for the Neural Basis of Cognition, University of Pittsburgh, Pittsburgh, Pennsylvania 15213

Visual attention, which improves perception of attended locations or objects, has long been known to affect many aspects of the responses of neuronal populations in visual cortex. There are two nonmutually exclusive hypotheses concerning the neuronal mechanisms that underlie these perceptual improvements. The first hypothesis, that attention improves the information encoded by a population of neurons in a particular cortical area, has considerable physiological support. The second hypothesis is that attention improves perception by selectively communicating relevant visual information. This idea has been tested primarily by measuring interactions between neurons on very short timescales, which are mathematically nearly independent of neuronal interactions on longer timescales. We tested the hypothesis that attention changes the way visual information is communicated between cortical areas on longer timescales by recording simultaneously from neurons in primary visual cortex (V1) and the middle temporal area (MT) in rhesus monkeys. We used two independent and complementary approaches. Our correlative experiment showed that attention increases the trial-to-trial response variability that is shared between the two areas. In our causal experiment, we electrically microstimulated V1 and found that attention increased the effect of stimulation on MT responses. Together, our results suggest that attention affects both the way visual stimuli are encoded within a cortical area and the extent to which visual information is communicated between areas on behaviorally relevant timescales.

Key words: attention; middle temporal area; primary visual cortex; variability

Significance Statement

Visual attention dramatically improves the perception of attended stimuli. Attention has long been thought to act by selecting relevant visual information for further processing. It has been hypothesized that this selection is accomplished by increasing communication between neurons that encode attended information in different cortical areas. We recorded simultaneously from neurons in primary visual cortex and the middle temporal area while rhesus monkeys performed an attention task. We found that attention increased shared variability between neurons in the two areas and that attention increased the effect of microstimulation in V1 on the firing rates of MT neurons. Our results provide support for the hypothesis that attention increases communication between neurons in different brain areas on behaviorally relevant timescales.

Introduction

Visual attention allows observers to concentrate on behaviorally relevant parts of a visual scene and, like other sensory, cognitive,

and motor processes, likely involves the coordinated activity of large groups of neurons in many brain areas. Attention has long been hypothesized to act by selecting the most important visual information for further processing, effectively filtering out information that is not behaviorally relevant. One possibility is that this selection is accomplished by increasing communication between neurons that encode the attended information in different cortical areas. This hypothesis has most frequently been tested on very fast timescales (e.g., a single millisecond, which may reflect synaptic efficacy; for review, see Womelsdorf and Fries, 2007). Longer timescales, however, may be extremely important behaviorally because, in natural vision, primates typically look around a visual scene, effectively taking a new snapshot of visual infor-

Received Feb. 24, 2016; revised June 2, 2016; accepted June 4, 2016.

Author contributions: D.A.R. and M.R.C. designed research; D.A.R. performed research; D.A.R. and M.R.C. contributed unpublished reagents/analytic tools; D.A.R. and M.R.C. analyzed data; D.A.R. and M.R.C. wrote the paper.

This work was supported by the National Institutes of Health (Grants 4R00EY020844-03 and R01 EY022930 to M.R.C., Training Grant 5T32NS7391-14 to D.A.R., and Core Grant P30 EY008098), the Whitehall Foundation (M.R.C.), a Klingenstein-Simons Fellowship (M.R.C.), the Simons Foundation (M.R.C.), a Sloan Research Fellowship (M.R.C.), and a McKnight Scholar Award (M.R.C.). We thank David Montez for assistance with animal training and recordings; Karen McCracken for technical assistance; and Joshua Alberts, Gabriela Costello, Adam Kohn, John Maunsell, and Amy Ni for comments on an earlier version of the manuscript.

The authors declare no competing financial interests.

Correspondence should be addressed to Douglas A. Ruff, Department of Neuroscience and Center for the Neural Basis of Cognition, University of Pittsburgh, 4400 Fifth Ave., MI Room 115, Pittsburgh, PA 15213. E-mail: ruffd@pitt.edu.

DOI:10.1523/JNEUROSCI.0610-16.2016
Copyright © 2016 the authors 0270-6474/16/367523-12\$15.00/0

mation, approximately every 200–300 ms. Changes in the extent to which the trial-to-trial variability in the responses of pairs of neurons is correlated (termed spike count correlations or r_{SC} , which are typically calculated using the number of spikes counted over hundreds of milliseconds) may be a signature of the communication of sensory information between areas and are mathematically nearly independent from measurements of millisecond-level synchrony (Bair et al., 2001).

We explored the hypothesis that visual attention increases interarea communication on longer timescales by recording simultaneously from multiple neurons with overlapping receptive fields in both primary visual cortex (V1) and the middle temporal area (MT). V1 and MT are heavily interconnected areas that both encode information about visual motion (Born and Bradley, 2005). During the recordings, two monkeys performed a motion direction change detection task that required them to shift attention between three moving stimuli. We found that, in contrast to its effects on spike count correlations between pairs of neurons in the same cortical area (Cohen and Maunsell, 2009, 2011; Mitchell et al., 2009; Zénon and Krauzlis, 2012; Herrero et al., 2013; Gregoriou et al., 2014; Ruff and Cohen, 2014a; Luo and Maunsell, 2015), attention increases spike count correlations between pairs of neurons in different areas. This result is qualitatively consistent with the idea that shifts in attention are associated with changes in the communication between neurons in different areas on the timescale of perceptual decisions.

To test the hypothesis that attention changes communication between V1 and MT in a complementary, causal manner, we measured the effect of electrically microstimulating V1 neurons on the activity of MT neurons in different attention conditions. We found that attention increases the extent to which V1 activity affects the activity of downstream neurons in MT, suggesting that attention increases the influence of V1 responses on MT responses.

Together, our results provide novel correlative and causal support for the hypothesis that attention affects the influence of the activity of neurons in one area on the responses of downstream neurons. More broadly, they provide evidence that attention is associated with at least two types of changes in visual cortex, affecting both the way visual stimuli are encoded within a cortical area and the extent to which visual neurons in different areas interact with each other.

Materials and Methods

Electrophysiological recordings. The subjects in this experiment were two adult male rhesus monkeys (*Macaca mulatta*, 8 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 behavioral training, we implanted each animal with a titanium head post. After the animal learned the task (4–5 months), we implanted a 10×10 microelectrode array (Blackrock Microsystems) in area V1 and a recording chamber that gave us access to area MT. The V1 array was connected to a percutaneous connector that allowed simultaneous recordings from 96 electrodes. The distance between adjacent electrodes on the array was $400 \mu\text{m}$ and each electrode was 1 mm long. We identified area V1 using stereotactic coordinates and by visually inspecting the sulci. During each recording session, we inserted either a single electrode (Fred Haer Corporation) or a 24-channel linear probe into MT. We identified MT using stereotactic coordinates, gray and white matter transitions, and receptive field properties.

We recorded neuronal activity from V1 and MT simultaneously during daily experimental sessions for several months in each animal. We searched for MT units with firing rates driven above baseline by a single stimulus placed at each of two locations that would allow for both successful task performance and overlap with the receptive fields of the

recorded V1 units. We picked the stimulus locations so that the MT units responded similarly to individual stimuli at each location (the average difference in MT responses to individual stimuli at each of the two stimulus locations was 9.7 sp/s, or fewer than two spikes per 200 ms stimulus presentation). We included recording sessions for analysis when the MT unit's receptive field largely overlapped the envelope of receptive fields of the units we recorded in V1 and when the animal completed at least 150 behavioral trials (mean = 648 completed trials from the 32 sessions used in Figure 3). For the between-area correlation analyses (see Figs. 3, 4, and 5), we recorded from a single electrode in MT and optimized the visual stimuli for the tuning properties of that unit ($n = 12$ experimental sessions from Monkey 1 and 20 experimental sessions from Monkey 2).

To measure correlations between pairs of MT units (see Fig. 2G; $n = 16$ experimental sessions from Monkey 1 and 15 experimental sessions from Monkey 2) and in the microstimulation experiments (see Fig. 6; $n = 16$ experimental sessions from Monkey 1), we recorded with 24-channel probes (either Alpha Omega Linear Micro Arrays with a contact spacing of $60 \mu\text{m}$ or Plexon V-Probes with a contact spacing of $50 \mu\text{m}$). In these experiments, it was impossible to optimize the visual stimuli for all MT units.

To determine whether we measured spikes from the same neuron on multiple electrodes of our 24-channel probes, we examined cross correlograms between all pairs of electrodes. It was clear that spikes recorded on adjacent electrodes were largely independent, but we did observe significantly more synchrony on adjacent channels than we did on more distant channels. Therefore, we only analyzed neuronal responses from nonadjacent channels, effectively discarding half of the data that we collected. Including data collected from all channels did not qualitatively change any results, including the effect of spatial attention on noise correlations between pairs of MT units (see Fig. 2G) or the effect of attention on altering MT firing rates in response to V1 microstimulation (see Fig. 6).

All spike sorting was done offline manually using Offline Sorter (version 3.3.2; Plexon). We based our analyses on both single units and multiunit clusters and use the term "unit" to refer to either. A previous study measuring spike-rate correlations from single and multiunits in V4 (Cohen and Maunsell, 2009) found that, whereas the magnitude of r_{SC} depended strongly on firing rate, the effects of attention did not. That is, attention seemed to modulate r_{SC} by a fixed proportion. Unfortunately, our current dataset does not contain enough single units to do similar analyses. Although it is almost certainly true that the preponderance of multiunits in our dataset affects the absolute magnitude of the correlation coefficients that we measured, we think it is extremely unlikely that it accounts for the effect of attention that we observed. For most analyses, we included a V1 unit if it responded significantly more to a full contrast stimulus at one of the two locations than the other, collapsed across the two directions of motion (t test, $p < 0.01$). We experimented with more stringent criteria for inclusion and found that they made qualitative, but not quantitative, differences in our results. For example, when we redid our analyses with a more stringent criterion (including only neurons with responses to the two stimuli that were significantly different at $p < 0.0001$ rather than $p < 0.01$), the magnitude of all of the attention effects increased. The mean attention-related changes in V1 rates (quantified as attention indices) and within-V1 spike count correlation increased as well (V1 attention indices, mean = 0.012, $p = 2.4 \times 10^{-49}$; within V1 r_{SC} for attend in = 0.184, attend out = 0.196, 25,636 pairs, $p = 9.7 \times 10^{-31}$; V1–MT r_{SC} attend in = 0.0194, attend out = -0.00053 , $p = 4.96 \times 10^{-7}$). For the analyses of attention-related changes in r_{SC} for spatial attention between pairs of MT units (see Fig. 2G), we included MT units if they responded significantly more to a full contrast stimulus at one of the two locations than the other, collapsed across stimulus direction.

Visual stimuli, eye position, and behavior. We presented visual stimuli using custom software (written in MATLAB using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) on a CRT monitor (calibrated to linearize intensity; 1024×768 pixels; 120 Hz refresh rate) placed 57 cm from the animal. We monitored eye position using an infrared eye tracker (Eyelink 1000; SR Research) and 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 (30,000 samples/s) using hardware from Ripple.

The monkeys performed a motion direction change detection task (see Fig. 1A). A trial began when the monkey fixated a small spot within a 1° square fixation window in the center of the video display. Monkeys typically maintained fixation within a much smaller window than was allotted (the median SD of eye position during analyzed stimulus presentations was 0.18°) and we did not include trials for analysis during which we detected a microsaccade within the fixation window (see “Data analysis” section).

In the blocks of trials used to measure attention, we presented two achromatic Gabor patches with size and location selected so that they both lay within the receptive field of the MT unit under study. The two stimuli were approximately equidistant from the fixation spot (their distances from the fixation spot never differed by $>0.3^\circ$). The stimuli were centered $\sim 2.5\text{--}3.5^\circ$ eccentric and each stimulus typically subtended $<1^\circ$ of visual angle. The visual stimuli flashed on for 200 ms and off for an interval that was randomly selected from a uniform distribution with a range of 200–400 ms. During the blocks of trials when the animal directed its attention toward one of the stimuli inside the receptive field of the MT neuron, the two stimuli were either both 8% contrast or both 100% contrast. Stimulus contrast was randomly interleaved on each stimulus presentation.

In separate blocks of trials (used to determine the tuning of V1 cells), the animal was instructed to direct its attention to a third stimulus in the opposite hemifield (attend opposite blocks). In these blocks of trials, the two stimuli in the receptive field were independently presented at 0%, 50%, or 100% contrast. The third, attended, stimulus was presented at either 8% or 100% contrast, with its contrast randomly interleaved on each stimulus presentation and independently selected from the contrasts of the stimuli in the MT receptive field.

In all trials, all of the stimuli drifted at the same speed, which was selected from a range between 6 and $12^\circ/\text{s}$ (picked to elicit large responses in the MT unit). The two stimuli in the MT receptive field moved in opposite directions (the preferred and null directions of the MT cell under study) and which of the two stimuli moved in the preferred direction varied randomly from trial to trial. The third stimulus, when present during attend opposite blocks, moved in an orthogonal direction to the stimuli in the receptive field.

After an unsigned number of stimulus presentations selected from an exponential distribution (minimum two stimulus presentations, mean six, maximum 14), the direction of one of the stimuli changed. During each experimental session, we selected a single magnitude of the direction change designed to get the animal to perform near psychophysical threshold (range $10\text{--}45^\circ$). The probability of the direction change was independent at each location (the unattended stimuli changed on $\sim 12\%$ of trials). Before the start of each block of trials, the monkey performed 5–10 instruction trials (which were not included in any analyses) in which there was only a single stimulus. The location of this stimulus constituted a cue as to the attended location. In the upcoming block of trials, if the attended stimulus was the one that changed, the monkey was given a liquid reward for making a saccade to that stimulus within 500 ms of the change. To account for saccadic latency and to avoid rewarding the monkey for guessing, the monkey was rewarded only for saccades that began at least 100 ms after the change. If no saccade occurred within 500 ms of the direction change at the cued location, then the trial was classified as a miss and terminated with no reward. If no change occurred within the maximum 14 stimulus presentations, the monkey was rewarded simply for maintaining fixation. Attention was cued to one stimulus location or the other in blocks of 50–100 trials. The monkey was never rewarded for making a saccade to distractor changes.

Overall, the monkeys correctly detected the stimulus changes on 66% of completed trials during the blocks of trials in which attention was directed to one of the stimuli inside the MT unit's receptive field across the 32 experimental sessions in which we collected neuronal data from both V1 and MT. During these blocks of trials, the monkeys detected 90% of the full-contrast direction changes and 42% of the low-contrast direction changes that occurred at the attended location. The monkeys responded to 21% of the unattended orientation changes, but were not rewarded for these responses. During the attend opposite blocks (which were much easier for the animals because the distractors were far away

from the attended stimuli), they detected direction changes at full contrast 99% of the time and responded to direction changes at the unattended location 5% of the time. Responses to individual stimuli during these blocks were used to determine which of the two stimuli within the MT unit's receptive field elicited the biggest response in each V1 unit and to compute the location and direction selectivity of each V1 unit (see Fig. 4).

Microstimulation experiments. A subset of our experiments involved measuring the number of extra spikes in MT that were produced by microstimulating V1 (see Fig. 6). We stimulated on the V1 array using 200 Hz biphasic pulses (pulse duration = $0.2\ \mu\text{s}$). We increased the current (range $30\text{--}65\ \mu\text{A}$) and number of electrodes (range 1–3 electrodes) until the V1 microstimulation produced a readily observable modulation of the responses of the MT units. If we could not elicit a response to V1 microstimulation in MT using this approach, we moved the probe in MT until we could. We stimulated for a 50 ms period during randomly interleaved stimulus presentations in the direction change discrimination task. Microstimulation never occurred on the first visual stimulus presentation and occurred with 50% probability on subsequent stimulus presentations. In different experiments, we began microstimulating at either 50 or 100 ms after stimulus onset (always for a 50 ms period). The effects of V1 microstimulation on MT responses did not qualitatively depend on when the stimulation began, so we combined all of the data in Figure 6C.

We saw evidence of electrical artifacts in our V1 recordings, but not in our MT recordings. In V1, we saw large artifacts despite fast-settle amplifiers and other hardware technologies (Ripple) designed to minimize artifacts from electrical stimulation. These artifacts were obvious because they were huge in voltage and were time locked to the stimulation pulses. For that reason, we only report findings from data collected in MT that were collected using a separate front-end amplifier from the one that was used to stimulate. Because we used separate (isolated) amplifiers, we were able to record in MT during stimulation without seeing evidence of the sort of artifacts that occurred in V1.

We took several additional steps to be sure that our MT results were not contaminated by more subtle artifacts. First, we listened to the MT channels while recording and stimulating and did not hear artifacts that sounded distinct from recordings in the absence of V1 microstimulation. Second, we frequently observed no change in MT responses when we stimulated in V1. We often had to stimulate on different V1 electrodes and move our electrodes in MT before we could find any observable change in MT responses. Third, we used the approach used by Moore and Armstrong (2003) of basing our analyses on the time period after microstimulation to eliminate the possibility of artifacts locked to the stimulation pulses. Perhaps most importantly, our analysis of MT responses during these experiments compared responses from different attention conditions that would effectively subtract out any possible (insidious) electrical/recording artifacts that could have resulted from microstimulation.

We designed this experiment so that changes in eye position or the inducement of microsaccades as a result of microstimulation would be unlikely to affect our data. We looked at responses in MT during a 100 ms period of time that followed a 50 ms microstimulation period (therefore, the longest point of time included was 150 ms from the onset of microstimulation); 150 ms would be very fast to both detect a phosphene and execute an eye movement toward it. Furthermore, longer latency eye movements of the type that might be in response to a phosphene, would take even longer to evoke changes in neuronal responses in visual cortex. Accordingly, microstimulation was not associated with any differences in the animals' eye velocity at any time during the visual stimulus presentation (comparison of eye velocity before microstimulation vs a period of time during and after microstimulation; Wilcoxon ranked-sum test, $p = 0.6$).

Data analysis. Our analyses were performed on responses to all 100% contrast stimulus presentations during correct trials with the exception of the first stimulus presentation (to reduce the impact of adaptation on our results) and only included stimulus presentations before the first direction change that occurred in a trial at any stimulus location. We excluded low-contrast stimulus presentations from our analyses of the effects of attention on neuronal responses because they typically elicited

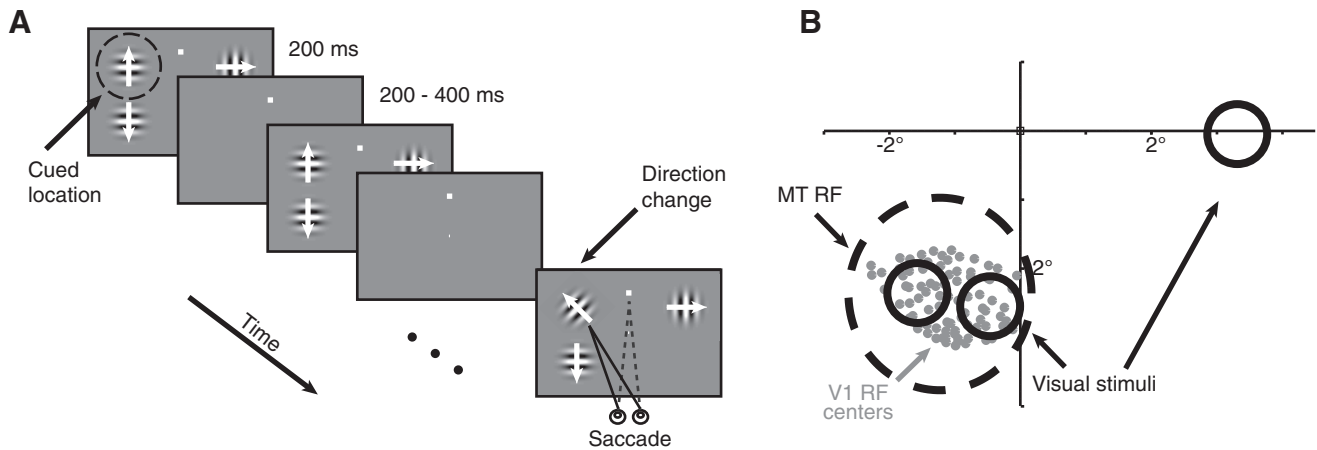


Figure 1. Psychophysical and physiological methods. **A**, Schematic of the motion direction change detection task. Once the monkey fixated a central spot, three small Gabor stimuli synchronously flashed on for 200 ms and off for a randomized interval between 200 and 400 ms. Two stimuli were positioned inside the joint receptive fields of the V1 and MT units that we recorded and a third was positioned in the opposite hemifield. After an unsigned and randomized number of stimulus presentations (picked from an exponential distribution, minimum two, mean six, maximum 14 stimulus presentations), one stimulus moved in a different direction. The monkeys were cued in blocks of 50–100 trials to detect changes in (and therefore attend to) one of the stimuli and ignore motion direction changes in the other stimuli. The monkey was rewarded for making a saccade to the attended stimulus within 500 ms of the stimulus change. Responses to distractor changes were never rewarded. The two stimuli in the MT receptive field moved in opposite directions (the preferred and null directions of the MT cell under study) and which of the two stimuli moved in the preferred direction varied randomly from trial to trial. The third stimulus, when present, moved in an orthogonal direction. **B**, Receptive fields and visual stimuli from an example experimental session. We recorded simultaneously from a 96-channel chronically implanted microelectrode array in area V1 and either a single electrode or a movable 24-channel probe in area MT. We selected MT units with receptive fields (black dashed circle) that overlapped the envelope of receptive fields of the units that we recorded in V1 (centers denoted by the gray dots). The locations and approximate sizes of the visual stimuli are denoted by solid black circles. Spatial receptive fields were estimated by rapidly presenting a single, small Gabor stimulus at a range of locations while the monkey fixated.

very low firing rates, which affect estimates of spike count correlations (for review, see Cohen and Kohn, 2011). Stimulus presentations in which a microsaccade was detected anywhere between 10 ms before until 10 ms after the stimuli were shown were excluded from analysis. We identified microsaccades using a velocity detection algorithm (Engbert and Kliegl, 2003) and detected them on an average of 6.1% of usable stimulus presentations across all experimental sessions (SEM = 0.4%). To determine the effects of attention on spike count correlations, we calculated correlations separately for each condition (i.e., for each of the two possible positions of the preferred and null direction stimuli) and then averaged across conditions for each pair.

Small deviations in eye position can affect measurements of r_{SC} because they can cause variability in neuronal responses. We addressed this issue in two ways. First, if variability in eye position had a large effect on response variability, we would expect the V1 neurons we recorded (because they have small receptive fields) to have unusually high trial-to-trial response variability. However, the Fano factors (ratio of the variance to the mean spike count response) of the V1 neurons we recorded were not unusually high (the average Fano factor was 1.22 ± 0.01 SEM and 1.24 ± 0.02 SEM when the monkey attended or did not attend the stimulus that best overlapped the V1 neuron's receptive field, respectively). Second, we extracted the slope of the best-fit line relating firing rate and vertical or horizontal eye position for each neuron (for both V1 and MT). Although r_{SC} trivially depends on the slope of those lines (even for pure noise), the attention-related change in r_{SC} does not depend on the relationship between eye position and firing rate. The mean attention-related change in r_{SC} for V1–MT pairs that have both slopes (horizontal and vertical) the same sign was 0.017 ($n = 409$ pairs), for pairs with one slope the same sign was 0.020 ($n = 539$ pairs), and for pairs with both slopes with opposite sign was 0.017 ($n = 683$ pairs), none of which are significantly different from each other (t test). This analysis implies that, whereas eye movements may certainly affect r_{SC} , they are not a likely source of the attention-related changes that we observed.

Using chronically implanted microarrays (as we used for our V1 recordings), it is difficult to tell whether we recorded from the same single or multiunit clusters on subsequent days. In an effort to deal with this concern, we analyzed the effects of attention on V1 units recorded during a single example recording session (see Figs. 2A–C, 3A). We selected this session because the animal performed a large number of trials with good

psychophysical performance and because recording quality was good. In addition, we summarized our entire V1 dataset (see Fig. 2D, E), but care should be taken in the interpretation of these data because the same units may contribute to multiple sessions. Importantly, the MT units were recorded on different electrodes each day, so each MT unit and V1–MT pair was unique.

The analyses in Figures 2, 3, and 4 are based on spike counts measured 30–230 ms after stimulus onset for V1 and 50–250 ms after stimulus onset for MT to account for the average visual latencies of neurons in both areas. Using identical windows for both areas led to qualitatively similar results to those presented here. The cross-correlograms (CCGs) in Figure 5 were trial-shuffle corrected to remove stimulus-induced correlations (Bair et al., 2001) and the population average was smoothed with a 5 ms Gaussian kernel. The analysis in Figure 6C is based on spike count responses calculated from the period of time 1–101 ms after the end of electrical microstimulation.

The distribution matching procedure to control for mean firing rate in our analysis of cross-area correlations was described in detail previously (Churchland et al., 2010; Ruff and Cohen, 2014b). Briefly, the goal of this analysis was to have the same distribution of geometric mean firing rates (but not covariances) in each of the four stimulus and attention conditions, so we used a different subdistribution of V1–MT pairs in each stimulus and attention condition. We compared distributions of means from each stimulus and attention condition and selected the greatest common distribution. We then subsampled our pairs of units in each condition to match that distribution and then analyzed spike count correlations for those subdistributions. There was a large overlap of these four firing rate distributions and 1267 pairs survived this mean matching. The mean-matched resampling process was repeated 1000 times and we reported the average r_{SC} values from these resampled distributions and the SEM from one representative resampled distribution.

Results

Attention increases rates and decreases correlations within each cortical area

We trained two rhesus monkeys to perform a motion direction change detection task (Fig. 1A) while we simultaneously recorded from a 96-channel chronically implanted microelectrode array in

V1 and either a single electrode or a movable 24-channel probe in MT. Our dataset included both single units and multiunits and we use the term “unit” to refer to either (see Materials and Methods). We selected MT units with receptive fields that overlapped the span of the receptive fields of the units that we recorded in V1 (Fig. 1*B*). During blocks of trials when the animal directed attention toward the receptive fields of the neurons we recorded, we placed two drifting Gabor stimuli within the receptive field of the MT unit; one stimulus moved in the MT unit’s preferred direction and the other moved in the null direction. The location of the preferred stimulus varied on randomly interleaved trials. The stimuli were flashed synchronously on for 200 ms and off for between 200 and 400 ms. After a variable number of stimulus presentations, the direction of one of the stimuli changed. The monkeys were cued in blocks of trials to respond to changes in (and therefore attend to) one of the two stimulus locations and ignore motion direction changes in the other stimulus.

We analyzed the neuronal effects of attention with respect to the receptive field of each V1 unit. To identify the location that best drove each V1 unit, during a separate block of trials, we directed the animal’s attention to a third stimulus in the opposite visual hemifield, well outside of the receptive fields of the neurons under study. During these blocks of trials (which we term “attend opposite” blocks), the animal was cued to respond to changes in this third stimulus and ignore changes at the other two stimulus locations. During attend opposite blocks, we recorded responses to individual Gabors placed in the two locations inside the MT neuron’s receptive field moving in either the MT neuron’s preferred or null direction, as well as several additional stimulus configurations (see Materials and Methods). For each V1 unit with a response that was significantly different to individual stimuli at each of the two locations, we designated blocks of trials when the monkey directed attention into the receptive field of that unit as “attend in” when the monkey attended the stimulus that elicited the biggest response from that unit and “attend out” when the monkey attended the other location. The term “attend out” is shorthand because a stimulus in the worse location often still elicited a response in the V1 unit. This aspect of our experimental design would not lead to a bias in the results that we report. Indeed, including V1 units that are well driven by both stimuli would only serve to weaken our reported effect size (see Fig. 4*A*).

Using this experimental design, we were able to replicate the previously reported hallmarks of attention-related changes in visual neurons. Attention directed into the receptive fields of visual neurons has been typically associated with higher firing rates (Maunsell and Cook, 2002; Yantis and Serences, 2003; Reynolds and Chelazzi, 2004) and lower spike count correlations with other similarly tuned neurons (Cohen and Maunsell, 2009, 2011; Mitchell et al., 2009; Zénon and Krauzlis, 2012; Herrero et al., 2013; Gregoriou et al., 2014; Ruff and Cohen, 2014a; Luo and Maunsell, 2015). Consistent with these results, we found that attention increased mean firing rates in V1 (quantified using a standard attention index; see Fig. 2*A,D*) and decreased spike count correlations, both in the same example recording session as the data depicted in Figure 1*B* (Fig. 2*A,B*, mean attention index for 62 V1 units in the example session = 0.008 or ~1.6% modulation, and mean attention-related correlation decrease = 0.008) and for the entire population (Figs. 2*D,E*; mean attention index for 2178 V1 units across 42 recording sessions = 0.011 or ~2% modulation and mean attention-related correlation decrease = 0.008). Attention directed to the preferred stimulus of the MT unit was associated with higher firing rates, both in the same

example recording session (Fig. 2*C*; attention index = 0.047 or ~10% modulation) and for the population of 32 MT units (Fig. 2*F*; mean attention index = 0.05 or ~10% modulation). These attention-related changes in rates and correlations were consistent for both animals (Monkey 1 mean V1 attention index = 0.015, decrease in within V1 r_{SC} = 0.014, mean MT attention index = 0.04; Monkey 2 mean V1 attention index = 0.007, decrease in within V1 r_{SC} = 0.004, mean MT attention index = 0.067; Wilcoxon signed-rank and rank-sum tests, all comparisons significant at $p < 0.001$). In a subset of experimental sessions, we recorded multiple MT units simultaneously. Because of our recording methods, these units had a broad range of spatial receptive fields. When two MT units had only partially overlapping receptive fields, we found that, consistent with our V1 results, attention to the stimulus in their joint receptive field was associated with lower correlations than attention directed elsewhere (Fig. 2*G* and Materials and Methods, mean spatial attention-related spike count correlation decrease = 0.019).

The attention-related changes that we observed in firing rates and in spike count correlations between pairs of neurons in the same cortical area are on the small end of those reported in the literature. Several factors, including task difficulty and the typically small cognitive effects in V1, likely contribute to the magnitude of the attention-related modulations that we observed (for review, see Maunsell and Cook, 2002). One factor unique to our study likely diminished the size of the changes that we observed. Because we needed to place our electrode array on the surface of the cortex, the receptive fields of the V1 units that we recorded were small and close to the fovea (and therefore the receptive fields of the MT units with overlapping receptive fields were also close to the fovea). To fit two stimuli within the MT neuron’s receptive field, we therefore had to use very small stimuli that were very close to each other. There may be limits on the extent to which animals can switch attention between two very small, nearby locations. Put another way, the “spotlight of attention” may be large enough that it necessarily included the unattended stimulus to a certain extent (Intriligator and Cavanagh, 2001). This factor likely affected both our V1 and MT results. Critically, despite the fact that we asked monkeys to switch attention between two small, nearby locations, our results were qualitatively similar to those in previous studies.

Attention has opposite effects on correlations within and between cortical areas

We next measured correlations between pairs of neurons in different cortical areas. These measurements allowed us to dissociate between two different hypotheses. If attention affects interactions between pairs of neurons without regard to which area the neurons are in, then attention toward a V1 neuron’s receptive field should decrease its spike count correlations with MT, just as it does for pairs within each area. Conversely, if attention increases communication between cortical areas, then attention toward a V1 neuron’s receptive field should increase, rather than decrease, the trial-to-trial variability that it shares with MT.

We found strong evidence for this second hypothesis. We measured consistently higher spike count correlations between pairs of V1 and MT units when attention was directed to their joint receptive field (Fig. 3*A* example recording session, 62 V1–MT pairs, mean attention-related difference = 0.052; Fig. 3*B* population average of 1631 pairs across 32 recording sessions, mean attention-related difference = 0.018). This attention-related increase in correlations between V1 and MT was significant for each animal (Monkey, mean r_{SC} increase with

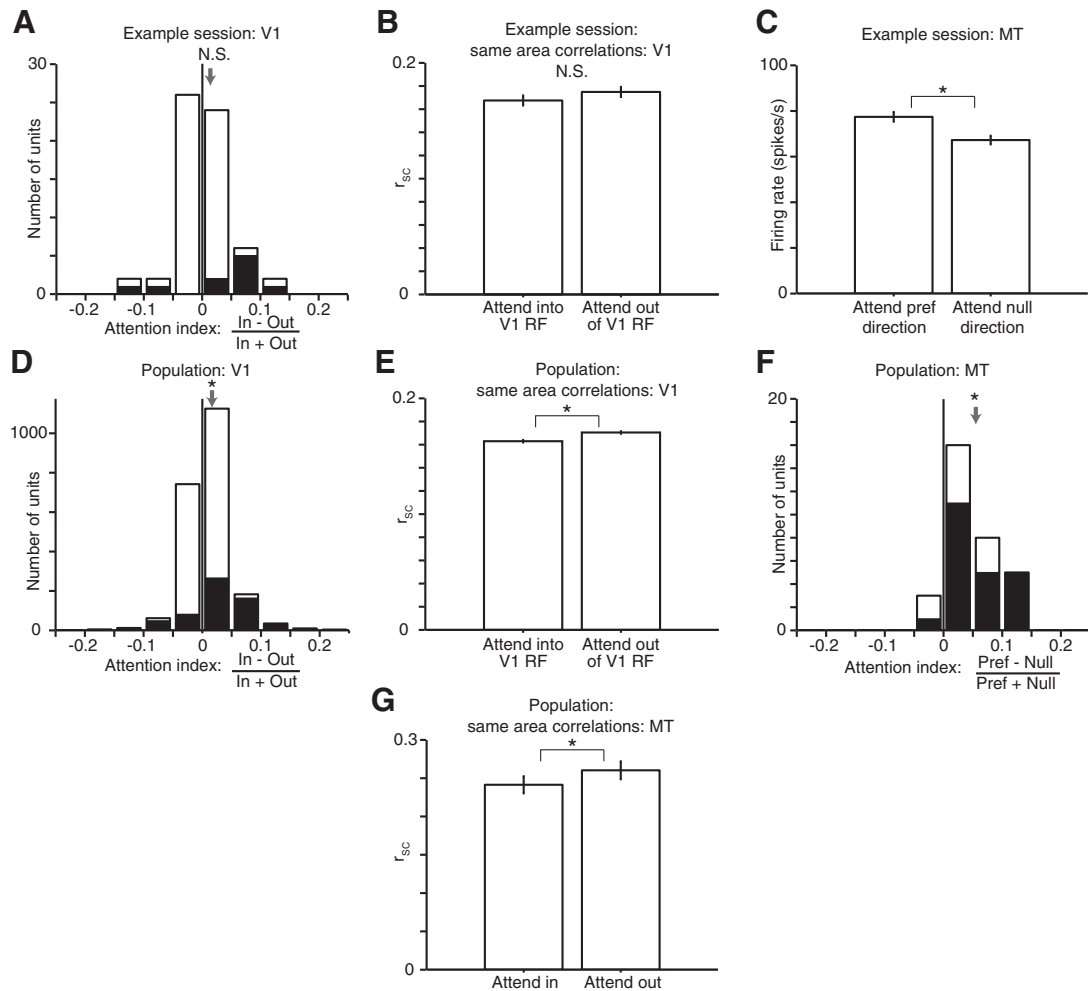


Figure 2. Attention increases rates and decreases correlations in both V1 and MT. **A**, Histogram of attention indices for the V1 units from the example experimental session for which receptive fields are plotted in Figure 1*B*. The attention index quantifies attention-related changes in mean firing rate as the difference divided by the sum of the firing rates in the two attention conditions (62 units in this example session; the mean attention index = 0.008 is indicated by the gray arrow, Wilcoxon signed-rank test, $p = 0.097$). Shaded bars indicate units with attention indices that differed from 0 (t tests, $p < 0.05$). **B**, Mean r_{sc} between pairs of V1 units for the same example session in each attention condition. Attending to the joint receptive field of a pair of V1 units tended to decrease its spike count correlation ($n = 1048$ pairs, Wilcoxon ranked-sum test, $p = 0.17$). Error bars are ± 1 SEM in all panels. **C**, Attention-related changes in firing rate for the MT unit recorded during the same example session. As was typical in our dataset, the firing rate was significantly higher when the animal attended to the stimulus moving at the unit's preferred direction (Wilcoxon rank-sum test, $p = 0.007$). **D**, Histogram of attention indices for 2178 V1 units across 42 recording sessions in two animals (mean attention index 0.011, Wilcoxon signed-rank test, $p = 5.5 \times 10^{-47}$). Conventions as in **A**. **E**, Mean r_{sc} between simultaneously recorded V1 units across all 34,404 pairs in 42 sessions in each attention condition. Attention significantly decreased r_{sc} (Wilcoxon rank-sum test, $p = 4.9 \times 10^{-6}$). **F**, Histogram of attention indices for 32 MT units recorded in 32 recording sessions (one unit per session; mean attention index 0.05, Wilcoxon signed-rank test, $p = 3.5 \times 10^{-6}$). Conventions as in **A** and **D**. **G**, Mean r_{sc} between simultaneously recorded MT units in each spatial attention condition. These units were recorded on multielectrode probes in largely separate sets of experiments ($n = 270$ total pairs from 16 experimental sessions in Monkey 1 and 15 experimental sessions in Monkey 2; difference between attention conditions is statistically significant; Wilcoxon signed-rank test, $p = 0.017$).

attention = 0.014; Monkey 2, mean r_{sc} increase with attention = 0.022; Wilcoxon rank-sum test, both significant at $p < 0.001$). We consistently observed this attention-related increase in spike count correlations in individual MT units in our study (the mean attention-related change in correlation for the population of MT units was significantly different from 0, Wilcoxon signed-rank test on the changes in r_{sc} between 32 MT units and simultaneously recorded V1 units, $p = 0.029$) and it was often detectable even in individual MT units (11 of 32 MT units showed significant attention-related changes in correlation with V1 units; Fig. 3*C*, t test, $p < 0.05$).

On the surface, the attention-related increase in cross-area correlations that we observed seems at odds with the attention-related decreases in correlations between pairs of V1 or pairs of MT units. The idea that attention results in two V1 neurons becoming more correlated with the same MT neuron but less cor-

related with each other seems like a paradox and would be mathematically impossible for very large correlation changes. Although differences in recording apparatus (and therefore electrical noise) make it tricky to compare within-area and cross-area correlations quantitatively, it is worth pointing out that, even when attention was directed to their joint receptive fields, correlations between units in the same cortical areas were substantially higher than between pairs when the two units were in different cortical areas. This difference likely reflects the known differences in shared connectivity between neurons in the same and in different cortical areas (Maunsell and Van Essen, 1983; Ungerleider and Desimone, 1986). This difference in correlation magnitude is also necessary for the opposite attention-related effects on correlations within and across areas to be mathematically possible (i.e., to maintain a positive definite covariance matrix). One possibility is that the magnitude of attention-related changes in corre-




Figure 3. Attention increases r_{sc} between V1 and MT. **A**, Average cross-area r_{sc} between V1 units and the MT unit from the same example experimental session in Figures 1B and 2A. We compared r_{sc} on trials with full-contrast visual stimuli when the animal switched attention between the two stimuli inside the MT neuron's receptive field. Error bars represent \pm SEM and the attention-related increase in r_{sc} was statistically significant (62 V1–MT pairs, Wilcoxon ranked-sum test, $p = 0.03$). **B**, Average cross-area r_{sc} between V1 units and the MT unit from all 1631 cross-area pairs across 32 recording sessions. Error bars represent \pm SEM and the attention-related increase in r_{sc} was statistically significant (Wilcoxon ranked-sum, $p = 1.4 \times 10^{-4}$). **C**, Histogram of average attention-related changes in cross-area r_{sc} between each of 32 MT units and simultaneously recorded V1 units. The mean attention-related increase in r_{sc} was significantly greater than 0 (mean = 0.018, Wilcoxon signed-rank test, $p = 0.029$). Shaded bars indicate recording sessions for which the average attention-related change in cross-area r_{sc} was significantly different from 0 (t test, $p < 0.05$).

lated variability is limited by the necessity of maintaining this relationship.

Importantly, we observed attention-related correlation increases regardless of whether the stimulus in the joint receptive field moved in the MT unit's preferred direction (the data in Fig. 3 are collapsed across stimulus conditions; see Materials and Methods). This robustness is important because it suggests that the attention-related changes in correlation cannot be attributed to attention-related differences in the firing rates of the V1 or MT units. Supporting this idea, we observed attention-related increases in correlation even in subsets of units with matched distributions of firing rates in the four attention and stimulus conditions (see Materials and Methods; mean attend in correlation for subsets of pairs with matched rates = 0.015, SEM = 0.005; mean attend out correlation = -0.0035 , SEM = 0.0049; Wilcoxon ranked-sum test, $p = 4.0 \times 10^{-4}$). There was also no relationship between the geometric mean firing rate of V1 neurons and the attention-related difference in r_{sc} between V1 and MT ($\rho = -0.018$, $p = 0.48$). It is also worth noting that the opposite sign of the effect of attention on correlations between neurons in the same or in different cortical areas removes a lot of concerns about possible artifacts. For example, if attention were associated with increases in correlations between pairs of neurons in the same cortical area, then the increase in cross-area correlations may have been a simple byproduct of those within-area changes. Together, our results show that, in contrast to the observed attention-related decreases in correlations between neurons in the same cortical area, attention increases the trial-to-trial variability that is shared between V1 and MT.

Attention-related changes in V1–MT correlations are related to V1 receptive field location, not direction selectivity

The V1 neurons that we recorded varied both in the extent to which they responded differently to visual stimuli at the two locations (which we term location selectivity) and the extent to which they responded differently to stimuli moving in the MT neuron's preferred versus null direction (direction selectivity). Although a minority of V1 neurons are direction selective (Schiller et al., 1976; De Valois et al., 1982), these V1 neurons are thought to project preferentially to MT (Movshon and New-

some, 1996). The attention-related increases in V1–MT correlations that we observed could be consistent with two possibilities: attention may affect communication selectively between direction-selective V1 neurons and MT or the attention-related changes in V1–MT correlations that we observed may more generally reflect spatial attention and would therefore depend on the location selectivity, not direction selectivity of the V1 units.

To test these hypotheses, we computed location and direction selectivity indices for each V1 unit. We defined the location selectivity index (LSI) as the response of the V1 unit to a single, full contrast stimulus at location 1 minus the response to a single stimulus at location 2 divided by their sum. Therefore, cells that respond equally to stimuli at both locations (either because they respond badly to both or well to both) will have indices near zero, cells that respond strongly to stimulus 1 and weakly to stimulus 2 have large positive indices, and cells that respond strongly to stimulus 2 and weakly to stimulus 1 have large negative indices.

We found that the attention-related change in V1–MT r_{sc} depended strongly on the LSI of the V1 unit. The y -axis of Figure 4A shows V1–MT r_{sc} when the monkey attended to location 1 minus V1–MT r_{sc} when the monkey attended to location 2. Spike count correlations between V1 and MT units with near zero LSIs are on average unchanged by attention and modulation of r_{sc} by attention grows with LSI. The data in Figure 4A are binned so that 10% of the 1631 V1–MT pairs that we recorded are in each bin. The correlation between the attention-related change in r_{sc} and the V1 neuron's LSI was 0.12, which was significantly greater than 0 ($p < 10^{-9}$). We also calculated LSIs for each MT unit and found no significant difference in attention-related changes in r_{sc} between V1 and MT pairs with either similar or different signed LSI (same sign LSI r_{sc} attention-related change = 0.021, SEM = 0.005; different sign LSI r_{sc} attention-related change = 0.015, SEM = 0.005; t test, $p = 0.35$).

We also computed a direction selectivity index (DSI) for each V1 unit (response to a single stimulus moving in the MT neuron's preferred direction minus the response to a single stimulus moving in the MT neuron's null direction divided by the sum). Overall, there was a much narrower range of DSIs than LSIs in the population of V1 units that we recorded (x -axis of Fig. 4B), largely because the vast majority of V1 units were not strongly

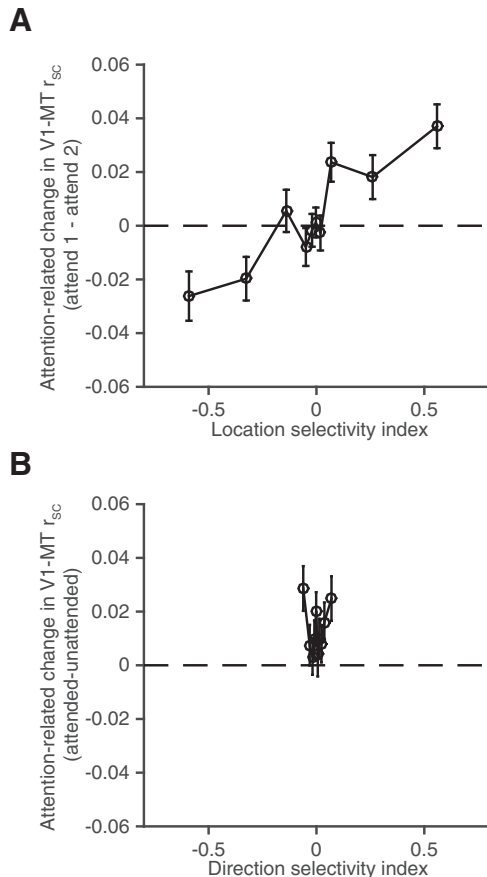


Figure 4. Attentional modulation of r_{SC} depends on location selectivity rather than on direction selectivity. **A**, Attention-related change in V1–MT r_{SC} (V1–MT r_{SC} when the monkey is attending to location 1 minus V1–MT r_{SC} when the monkey is attending to location 2) as a function of LSI for each V1 unit. The LSI is defined as the response of the V1 unit to a single, full contrast stimulus in location 1 minus the response to a single stimulus in location 2 divided by the sum. The data are binned such that 10% of V1–MT pairs fell in each bin. Error bars indicate SEM. **B**, Attention-related change in V1–MT r_{SC} (V1–MT r_{SC} when the monkey attended to the location that best drives the V1 cell minus V1–MT r_{SC} in the opposite attention condition) as a function of the DSI for the V1 unit. DSI is defined as the V1 unit's response to a single stimulus moving in the MT neuron's preferred direction minus the response to a single stimulus moving in the MT neuron's null direction divided by the sum. Conventions as in **A**.

direction selective. To give us the best chance to see any relationship between attention-related changes in r_{SC} and DSI, the y -axis in Figure 4B is V1–MT r_{SC} when the monkey attended to the location that best drives the V1 cell (even if that difference is not significant) minus the opposite attention condition. Even so, the attention-related change in V1–MT r_{SC} was uncorrelated with DSI (correlation between the attention-related change in r_{SC} and the V1 neuron's DSI was 0.01, which is indistinguishable from 0; $p = 0.60$).

Attention-related changes V1–MT r_{SC} have broad dynamics

We focused our correlation analyses primarily on the timescale of stimulus presentations (200 ms), but the spike count correlation measure that we used could capture changes at a range of time-scales. To investigate the dynamics of the attention-related correlation changes that we observed, we computed CCGs between V1 and MT for each attention condition (Fig. 5; see Materials and Methods). The CCGs were both relatively symmetric and broad and even the average across many pairs failed to show a sharp peak. This suggests that the vast majority of V1–MT pairs that we

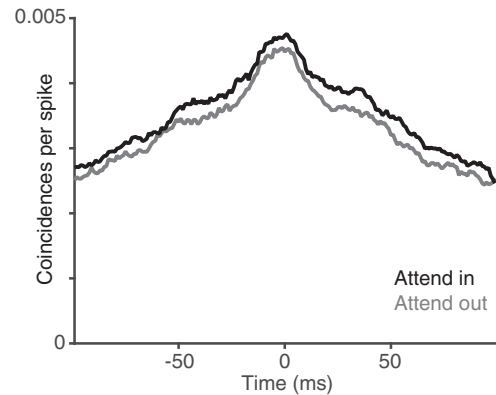


Figure 5. Attention affects correlations between V1 and MT consistently across time. Mean shuffle corrected CCGs between V1 and MT when attention was directed to (black) or away from (gray) the stimulus that best drove the V1 unit.

recorded were not connected by a single synapse, which is not surprising given that the probability of observing direct connections is very low (Movshon and Newsome, 1996; Zandvakili and Kohn, 2015). Therefore, the spike count correlations that we computed (which are related to the integral under the CCG; Bair et al., 2001) likely reflect co-fluctuations between populations of V1 and MT neurons rather than direct connectivity. Furthermore, the attention-related difference in the average CCGs was also broad and symmetric (the black and gray lines in Fig. 5 are approximately parallel), suggesting that the attention-related correlation changes have broad dynamics. To further investigate dynamics of attention-related correlation changes, we computed joint peristimulus time histograms (Aertsen et al., 1989) and also computed r_{SC} using sliding windows ranging from 5 to 200 ms. None of these analyses revealed interesting dynamics of correlation changes, suggesting that the attention-related change in r_{SC} is relatively constant across the 200 ms window that we used. However, further studies in which the stimuli are optimized for the recorded neurons will be necessary to thoroughly examine these issues.

Attention increases the extent to which manipulating V1 activity elicits extra spikes in MT

The observation that attention increases the correlation between activity of pairs of units in V1 and MT is broadly consistent with the hypothesis that attention increases communication between V1 and MT. To test this hypothesis in a complementary way, we performed an additional set of experiments using a causal manipulation. We hypothesized that, if neurons in V1 and MT communicate more when attention is directed toward their joint receptive fields, then attention would cause extra spikes in V1 to be reflected in MT responses more faithfully. This is the converse of the logic of the subthreshold stimulation protocol used to determine whether electrically stimulating frontal cortex could mimic the effects of attention on visual cortex (Moore and Armstrong, 2003).

We performed these experiments in the one animal with a V1 microelectrode array that lasted long enough for us to perform additional experiments after the correlation experiments were complete. We electrically stimulated electrodes on the microelectrode array in V1 using 50 ms trains of stimulation pulses (200 Hz, 30–65 μ A) during randomly interleaved stimulus presentations (on average, 50% of the stimulus presentations preceding the direction change) during the direction change detection task

(Fig. 6A; see Materials and Methods). Importantly, the stimulation occurred on subsets of randomly selected stimulus presentations within each trial rather than on separate groups of trials. Because stimulated and unstimulated stimulus presentations occurred within a few hundred milliseconds of each other and within the same trial, this design minimizes the probability that stimulation alters the monkey's attentional state or affects behavior. Accordingly, microstimulation did not interfere with the monkey's frequency or location of early eye movements (false alarms). When attention was directed to the location of the receptive fields in V1 that received stimulation, the frequency of early saccades to that location was not affected by the occurrence of microstimulation (mean frequency of early saccades on stimulus presentations with no microstimulation = 1.5% of trials, SEM = 0.28%; mean frequency of early saccades on stimulus presentations with microstimulation = 1.1%, SEM = 0.29%; *t* test, $p = 0.25$). Similarly, the frequency of early saccades to the microstimulated location when attention was directed away from that location was not affected by the occurrence of microstimulation (mean frequency of early saccades on stimulus presentations with no microstimulation = 0.8% of trials, SEM = 0.25%; mean frequency of early saccades on stimulus presentations with microstimulation = 1.1%, SEM = 0.22%; *t* test, $p = 0.25$). Although the stimulation currents that we used were likely large enough that a monkey could be trained to detect them (Murphy and Maunsell, 2007; Dagnino et al., 2015), the statistics of our task (changes in the visual stimulus typically occurred after several stimulus presentations) required the monkeys to adopt a very conservative criterion for a behavioral response. The electrical stimulation was not predictive of changes in the visual stimulus and was therefore not behaviorally relevant. We suspect that the animals simply ignored the electrical stimulation, which explains the lack of an effect of electrical stimulation on false alarm rates.

The microstimulation in V1, however, was large enough to produce readily observable changes in the responses of simultaneously recorded MT units (Fig. 6B). V1 microstimulation led to heterogeneous changes in MT responses during the 50 ms microstimulation epoch. Microstimulating V1 caused significant increases in the rates of 20 units (Wilcoxon ranked-sum test, $p = 0.01$) relative to their responses during the same period during visual stimulus presentations when there was no microstimulation. This heterogeneity may have multiple sources, including the fact that we likely recorded from different cell types and cortical layers in MT. In addition, it has been observed that electrical stimulation activates fibers of passage near the electrode tip (Histed et al., 2009), which is likely to activate a heterogeneous population of cells. The heterogeneous modulation of MT firing rates during the period of microstimulation in V1 did not correlate with the attentional modulation of firing rates after microstimulation that will be discussed below ($\rho = 0.043$, $p = 0.57$).

We reasoned that, if attention increases interactions between cortical areas, then microstimulating V1 should have a larger effect on the responses of MT units when attention was directed toward the receptive fields of the stimulated neurons. We did not observe evidence of electrical artifacts from V1 microstimulation in our MT recordings (see Materials and Methods). To be conservative, we adopted the convention in previous studies (Moore and Armstrong, 2003) of measuring the microstimulation-related differences in MT responses during the period of time immediately after the electrical stimulation in V1 was over. We calculated the difference between MT responses in the 100 ms period immediately after electrical microstimulation in V1 and in

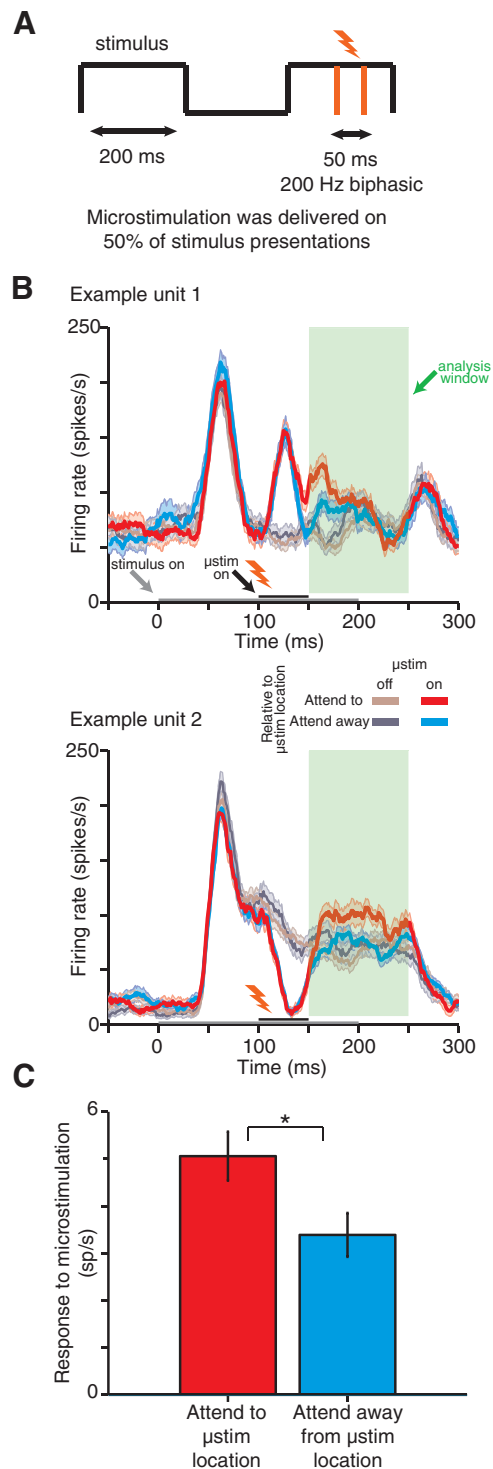


Figure 6. Attention increases the effect of manipulating V1 activity on the activity of MT units. **A**, Microstimulation was delivered for 50 ms pulses at 200 Hz during a randomly selected 50% of stimulus presentations in each trial. **B**, Example peristimulus time histograms from two example MT units. Attention increased the efficacy of microstimulation during the 100 ms period after V1 microstimulation (green shaded period). **C**, Average microstimulation efficacy (difference in firing rate during the 100 ms period shaded in **B** between stimulated and unstimulated visual stimulus presentations) for each attention condition for 183 units in 16 recording sessions. The difference in the response to microstimulation was significantly higher when the animal attended to the stimulus that overlapped the receptive fields of the microstimulated V1 units (mean increase, 1.67 sp/s, Wilcoxon signed-rank test, $p = 0.008$).

the analogous period during stimulus presentations in which there was no electrical stimulation. Attention increased this response difference, meaning that attention increased the number of extra MT spikes produced by microstimulation in V1 (Fig. 6C, mean increase with attention = 1.67 spikes/s or a 2.0% modulation of firing rate; Wilcoxon signed-rank test, $p = 0.008$). There was a similar attention-related difference in the extra MT spikes produced by V1 microstimulation if the analysis window also included the 50 ms microstimulation period: mean increase = 1.44 spikes/s. On an MT unit by unit basis, this attention-related increase was not significantly correlated with the raw efficacy of V1 microstimulation (the average difference between stimulated and unstimulated stimulus presentations, regardless of attention; correlation = 0.12, $p = 0.11$). Across the population, attentional modulation of the effect of microstimulation was correlated with the average firing rate of MT units ($\rho = 0.14$, $p < 0.05$). The attention-related increase in the effect of V1 microstimulation on MT responses, like the attention-related increase in spike count correlations between V1 and MT, is consistent with the hypothesis that attention increases communication between V1 and MT on long timescales.

Discussion

Implications for the neuronal mechanisms underlying attention

There are at least two nonmutually exclusive hypotheses about how attention improves perception. First, attention-related decreases in r_{SC} between neurons in the same visual area suggest that attention affects the way that visual stimuli are represented within a cortical area (Cohen and Maunsell, 2009, 2011; Mitchell et al., 2009; Zénon and Krauzlis, 2012; Herrero et al., 2013; Gregoriou et al., 2014; Ruff and Cohen, 2014a; Luo and Maunsell, 2015). Although the relationship between correlations and population coding is complicated (Shadlen et al., 1996; Abbott and Dayan, 1999; Ecker et al., 2011; Moreno-Bote et al., 2014), the effects of attention on correlations between pairs of neurons within a cortical area are broadly consistent with the hypothesis that attention improves perception by improving the fidelity with which visual stimuli are encoded.

There is another longstanding hypothesis that attention changes communication between cortical areas (for review, see Womelsdorf and Fries, 2007). This theory arose from observations that attention is associated with increased synchrony within and between sensory and higher-level cortical areas on fast (millisecond-level) timescales, which may reflect increased synaptic efficacy (Fries et al., 2001; Bichot et al., 2005; Womelsdorf et al., 2006; Saalmann et al., 2007; Lakatos et al., 2008; Gregoriou et al., 2009). A recent study also showed that attention increases synchrony between the thalamus and primary visual cortex on fast timescales, which likely reflects improved synaptic efficacy (Briggs et al., 2013).

We showed that attention increases interactions between visual cortical areas on longer timescales as well. This result is qualitatively consistent with recent work showing attention-related increases in longer timescale correlations between visual cortex and frontal cortex (Pooresmaeili et al., 2014) and between anterior cingulate cortex and frontal cortex (Oemisch et al., 2015). Changes in longer timescale correlations are not inconsistent with, but cannot be explained by, changes in synchrony because spike count correlations are dominated by variability on the timescale of tens to hundreds of milliseconds and are nearly independent from millisecond-level synchrony (Bair et al., 2001). Furthermore, we measured the extra spikes

in MT produced by V1 microstimulation by counting spikes during a 100 ms epoch after the microstimulation, which would not be affected by changes in precise synchrony.

We have phrased the logic of the microstimulation experiment using bottom-up wording (e.g., measuring the effect of increasing the activity of V1 neurons on MT responses). However, our results are also broadly consistent with the idea that attention changes interactions between V1 and MT in a top-down way or via a third area. Our results simply imply that attention increases functional communication between the two areas.

One potential question concerns how increases in correlation across areas might contribute to improvements in performance, especially because the attention-related r_{SC} increases were not restricted to direction-selective V1 units, which are thought to project selectively to MT (Movshon and Newsome, 1996). Correlated activity within V1 makes it possible that attention increases r_{SC} between MT and most V1 units, even if communication is preferentially increased between MT and the V1 units that project to MT. Such preferential increases in communication would likely increase r_{SC} between the MT neuron and all V1 neurons with responses that are more correlated with the projection neurons in the attended pool than with neurons in the unattended pool. Correlations within V1 decrease slowly with cortical distance or receptive field location (Kohn and Smith, 2005), so r_{SC} would likely increase between the MT neuron and all V1 neurons with receptive fields that overlap the attended stimulus. In the future, it will be interesting to determine whether, like correlation changes within a cortical area (Ruff and Cohen, 2014a), attention-related changes in cross-area correlations depend on the role that neurons play in a perceptual task.

Together, our results suggest that attention acts on visual cortex both by affecting the representation of visual stimuli within a cortical area and by increasing interactions between neurons in different areas. The relative importance of each of these mechanisms in improving perception will be an important avenue for further study.

Implications for models of how sensory information is combined across cortical areas

Our results also have implications for how sensory information is combined across cortical areas to guide decisions. Neurons in multiple visual cortical areas typically respond to any given visual stimulus and carry useful information for solving a particular perceptual task. For example, both V1 and MT carry information that would be useful for solving the motion direction change detection task that we used here. It is unknown, however, whether information from multiple areas is combined without regard to which area the neurons are in or whether the visual system processes information in a more sequential manner (e.g., with neurons from downstream areas reading out information from earlier processing stages).

The question of how sensory information is combined across cortical areas has been notoriously difficult to address. Simply identifying neurons in multiple areas that respond to a visual stimulus and carry information that might be useful for solving a particular perceptual task does not establish their role in the perceptual decision. Even showing that trial-to-trial variability in the activity of neurons in multiple areas is correlated with the animal's decision leaves open the possibility that neurons in one or both areas carry information that is simply correlated with, but not causally related to, the neurons that actually influence behavior (Nienborg and Cumming, 2010; Nienborg et al., 2012).

Our results suggest that there is something special about the boundaries between cortical areas because attention affects interactions between cortical areas in a qualitatively different way than interactions between neurons within the same area. Had we found that attention decreases spike count correlations between pairs of neurons in opposite areas just as it does for pairs within the same cortical area, we might have concluded that sensory information is simply combined from all neurons encoding task-relevant information. In general, our results suggest that measuring the way that cognitive factors such as attention affect interactions between different subgroups of neurons can be a useful tool for studying population coding.

Causal manipulations validate the use of spike count correlations as a probe for underlying cortical circuitry

The consistency of the results of our correlative and causal experiments validates the use of spike count correlations to probe cortical circuitry. Many studies over the past two decades have used spike count correlations to infer the structure and function of cortical circuits. For example, correlations have been used to infer functional connectivity (Aertsen et al., 1989; Alonso and Martinez, 1998) and information about cortical computations that is unavailable from the responses of individual neurons (for review, see Cohen and Kohn, 2011).

Although positive spike count correlations are widely thought to reflect common inputs to a pair of neurons, the relationship between correlations and connectivity is complicated. For example, a theoretical study showed that it is possible to have very low correlations in a densely interconnected network (Renart et al., 2010). A recent study showed that pairs of neurons in mouse visual cortex that have more common excitatory inputs tend to have higher trial-to-trial correlations in calcium responses (Ko et al., 2011). However, the idea that increases in spike count correlations reflect increased functional connectivity has been difficult to test experimentally.

Electrical microstimulation has long been used to assess functional connectivity in cortical circuits. However, these experiments are typically very difficult and they focus on timescales that are either much faster (Movshon and Newsome, 1996; Sommer and Wurtz, 2002; Berman and Wurtz, 2010) or much slower (Ekstrom et al., 2008; Moeller et al., 2008) than the timescale of perceptual decisions. Furthermore, to our knowledge, no other study has compared spike count correlations and the effects of electrical stimulation directly under identical conditions.

The consistency of our correlation and microstimulation results suggests that spike count correlations reflect important processes in the cortical circuit. Under both conditions in which attention increased spike count correlations between areas (when either the preferred or the null stimulus was in the V1 unit's receptive field (Ruff and Cohen, 2016), attention also increased the number of extra spikes in MT that were produced by electrical microstimulation in V1. These results suggest that, at least in our experimental conditions, increases in correlations correspond to increases in functional communication between two populations of neurons. As the technology for causally manipulating groups of neurons becomes more sophisticated and widespread, it will be interesting to determine the extent to which correlation changes correspond to changes in functional connectivity in different experimental systems and behavioral paradigms.

Neural computations have long been known to involve the activity of neurons throughout the brain. In this age of rapid advances in experimental technologies, recording from and manipulating the activity of large groups of neurons is more feasible

than ever before and combining these methods with models to place the results in context will be a powerful way to understand underlying neuronal computations. Together, our results suggest that measuring interactions between cortical areas will be critical to understanding the neuronal mechanisms underlying perception, cognition, and behavior.

References

- Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101. [CrossRef Medline](#)
- Aertsen AM, Gerstein GL, Habib MK, Palm G (1989) Dynamics of neuronal firing correlation: modulation of “effective connectivity.” *J Neurophysiol* 61:900–917. [Medline](#)
- Alonso JM, Martinez LM (1998) Functional connectivity between simple cells and complex cells in cat striate cortex. *Nat Neurosci* 1:395–403. [CrossRef Medline](#)
- Bair W, Zohary E, Newsome WT (2001) Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J Neurosci* 21:1676–1697. [Medline](#)
- Berman RA, Wurtz RH (2010) Functional identification of a pulvinar path from superior colliculus to cortical area MT. *J Neurosci* 30:6342–6354. [CrossRef Medline](#)
- Bichot NP, Rossi AF, Desimone R (2005) Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* 308:529–534. [CrossRef Medline](#)
- Born RT, Bradley DC (2005) Structure and function of visual area MT. *Annu Rev Neurosci* 28:157–189. [CrossRef Medline](#)
- Brainard DH (1997) The Psychophysics Toolbox. *Spat Vis* 10:433–436. [CrossRef Medline](#)
- Briggs F, Mangun GR, Usrey WM (2013) Attention enhances synaptic efficacy and the signal-to-noise ratio in neural circuits. *Nature* 499:476–480. [CrossRef Medline](#)
- Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, Bradley DC, Smith MA, Kohn A, Movshon JA, Armstrong KM, Moore T, Chang SW, Snyder LH, Lisberger SG, Priebe NJ, et al. (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci* 13:369–378. [CrossRef Medline](#)
- Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nat Neurosci* 14:811–819. [CrossRef Medline](#)
- Cohen MR, Maunsell JH (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* 12:1594–1600. [CrossRef Medline](#)
- Cohen MR, Maunsell JH (2011) Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron* 70:1192–1204. [CrossRef Medline](#)
- Dagnino B, Gariel-Mathis MA, Roelfsema PR (2015) Microstimulation of area V4 has little effect on spatial attention and on perception of phosphenes evoked in area V1. *J Neurophysiol* 113:730–739. [CrossRef Medline](#)
- De Valois RL, Yund EW, Hepler N (1982) The orientation and direction selectivity of cells in macaque visual cortex. *Vision Res* 22:531–544. [CrossRef Medline](#)
- Ecker AS, Berens P, Tolias AS, Bethge M (2011) The effect of noise correlations in populations of diversely tuned neurons. *J Neurosci* 31:14272–14283. [CrossRef Medline](#)
- Ekstrom LB, Roelfsema PR, Arsenault JT, Bonmassar G, Vanduffel W (2008) Bottom-Up Dependent Gating Visual Cortex. *Science* 321:414–417. [CrossRef Medline](#)
- Engbert R, Kliegl R (2003) Microsaccades uncover the orientation of covert attention. *Vision Res* 43:1035–1045. [CrossRef Medline](#)
- Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291:1560–1563. [CrossRef Medline](#)
- Gregoriou GG, Gotts SJ, Zhou H, Desimone R (2009) High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science* 324:1207–1210. [CrossRef Medline](#)
- Gregoriou GG, Rossi AF, Ungerleider LG, Desimone R (2014) Lesions of prefrontal cortex reduce attentional modulation of neuronal responses and synchrony in V4. *Nat Neurosci* 17:1003–1011. [CrossRef Medline](#)
- Herrero JL, Gieselmann MA, Sanayei M, Thiele A (2013) Attention-induced

- variance and noise correlation reduction in macaque V1 is mediated by NMDA receptors. *Neuron* 78:729–739. [CrossRef Medline](#)
- Histed MH, Bonin V, Reid RC (2009) Direct activation of sparse, distributed populations of cortical neurons by electrical microstimulation. *Neuron* 63:508–522. [CrossRef Medline](#)
- Intriligator J, Cavanagh P (2001) The spatial resolution of visual attention. *Cogn Psychol* 43:171–216. [CrossRef Medline](#)
- Ko H, Hofer SB, Pichler B, Buchanan KA, Sjöström PJ, Mrcic-Flogel TD (2011) Functional specificity of local synaptic connections in neocortical networks. *Nature* 473:87–91. [CrossRef Medline](#)
- Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci* 25:3661–3673. [CrossRef Medline](#)
- Lakatos P, Karmos G, Mehta AD, Ulbert I, Schroeder CE (2008) Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science* 320:110–113. [CrossRef Medline](#)
- Luo TZ, Maunsell JH (2015) Neuronal modulations in visual cortex are associated with only one of multiple components of attention. *Neuron* 86:1182–1188. [CrossRef Medline](#)
- Maunsell JH, Cook EP (2002) The role of attention in visual processing. *Philos Trans R Soc Lond B Biol Sci* 357:1063–1072. [CrossRef Medline](#)
- Maunsell JH, Van Essen DC (1983) The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. *J Neurosci* 3:2563–2586. [Medline](#)
- Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63:879–888. [CrossRef Medline](#)
- Moeller S, Freiwald WA, Tsao DY (2008) Patches with links: a unified system for processing faces in the macaque temporal lobe. *Science* 320:1355–1359. [CrossRef Medline](#)
- Moore T, Armstrong KM (2003) Selective gating of visual signals by microstimulation of frontal cortex. *Nature* 421:370–373. [CrossRef Medline](#)
- Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A (2014) Information-limiting correlations. *Nat Neurosci* 17:1410–1417. [CrossRef Medline](#)
- Movshon JA, Newsome WT (1996) Visual response properties of striate cortical neurons projecting to area MT in macaque monkeys. *J Neurosci* 16:7733–7741. [Medline](#)
- Murphey DK, Maunsell JH (2007) Behavioral detection of electrical microstimulation in different cortical visual areas. *Curr Biol* 17:862–867. [CrossRef Medline](#)
- Nienborg H, Cumming B (2010) Correlations between the activity of sensory neurons and behavior: how much do they tell us about a neuron's causality? *Curr Opin Neurobiol* 20:376–381. [CrossRef Medline](#)
- Nienborg H, Cohen MR, Cumming BG (2012) Decision-related activity in sensory neurons: correlations among neurons and with behavior. *Annu Rev Neurosci* 35:463–483. [CrossRef Medline](#)
- Oemisch M, Westendorff S, Everling S, Womelsdorf T (2015) Interareal spike-train correlations of anterior cingulate and dorsal prefrontal cortex during attention shifts. *J Neurosci* 35:13076–13089. [CrossRef Medline](#)
- Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis* 10:437–442. [CrossRef Medline](#)
- Pooresmaeili A, Poort J, Roelfsema PR (2014) Simultaneous selection by object-based attention in visual and frontal cortex. *Proc Natl Acad Sci U S A* 111:6467–6472. [CrossRef Medline](#)
- Renart A, de la Rocha J, Bartho P, Hollender L, Parga N, Reyes A, Harris KD (2010) The asynchronous state in cortical circuits. *Science* 327:587–590. [CrossRef Medline](#)
- Reynolds JH, Chelazzi L (2004) Attentional modulation of visual processing. *Annu Rev Neurosci* 27:611–647. [CrossRef Medline](#)
- Ruff DA, Cohen MR (2014a) Attention can either increase or decrease spike count correlations in visual cortex. *Nat Neurosci* 17:1591–1597. [CrossRef Medline](#)
- Ruff DA, Cohen MR (2014b) Global cognitive factors modulate correlated response variability between V4 neurons. *J Neurosci* 34:16408–16416. [CrossRef Medline](#)
- Ruff DA, Cohen MR (2016) Stimulus dependence of correlated variability across cortical areas. *J Neurosci* 36:7546–7556. [CrossRef](#)
- Saalmann Y, Pigarev IN, Vidyasagar TR (2007) Neural mechanisms of visual attention: how top-down feedback highlights relevant locations. *Science* 316:1612–1615. [CrossRef Medline](#)
- Schiller PH, Finlay BL, Volman SF (1976) Quantitative studies of single-cell properties in monkey striate cortex. I. Spatiotemporal organization of receptive fields. *J Neurophysiol* 39:1288–1319. [Medline](#)
- Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci* 16:1486–1510. [Medline](#)
- Sommer MA, Wurtz RH (2002) A pathway in primate brain for internal monitoring of movements. *Science* 296:1480–1482. [CrossRef Medline](#)
- Ungerleider LG, Desimone R (1986) Cortical connections of visual area MT in the macaque. *J Comp Neurol* 248:190–222.
- Womelsdorf T, Fries P (2007) The role of neuronal synchronization in selective attention. *Curr Opin Neurobiol* 17:154–160. [CrossRef Medline](#)
- Womelsdorf T, Fries P, Mitra PP, Desimone R (2006) Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* 439:733–736. [CrossRef Medline](#)
- Yantis S, Serences JT (2003) Cortical mechanisms of space-based and object-based attentional control. *Curr Opin Neurobiol* 13:187–193. [CrossRef Medline](#)
- Zandvakili A, Kohn A (2015) Coordinated neuronal activity enhances corticocortical communication. *Neuron* 87:827–839. [CrossRef Medline](#)
- Zénon A, Krauzlis RJ (2012) Attention deficits without cortical neuronal deficits. *Nature* 489:434–437. [CrossRef Medline](#)