

Using Neuronal Populations to Study the Mechanisms Underlying Spatial and Feature Attention

Marlene R. Cohen^{1,*} and John H.R. Maunsell¹

¹Harvard Medical School Department of Neurobiology and Howard Hughes Medical Institute, Boston, MA 02115, USA

*Correspondence: cohenm@pitt.edu

DOI 10.1016/j.neuron.2011.04.029

SUMMARY

Visual attention affects both perception and neuronal responses. Whether the same neuronal mechanisms mediate spatial attention, which improves perception of attended locations, and nonspatial forms of attention has been a subject of considerable debate. Spatial and feature attention have similar effects on individual neurons. Because visual cortex is retinotopically organized, however, spatial attention can comodulate local neuronal populations, whereas feature attention generally requires more selective modulation. We compared the effects of feature and spatial attention on local and spatially separated populations by recording simultaneously from dozens of neurons in both hemispheres of V4. Feature and spatial attention affect the activity of local populations similarly, modulating both firing rates and correlations between pairs of nearby neurons. However, whereas spatial attention appears to act on local populations, feature attention is coordinated across hemispheres. Our results are consistent with a unified attentional mechanism that can modulate the responses of arbitrary subgroups of neurons.

INTRODUCTION

Visual attention allows observers to focus on a subset of a complex visual scene. Spatial attention, which improves perception of stimuli at attended locations, has been well studied. However, observers can attend to many other attributes of a visual scene (Wolfe et al., 2004), including features (Haenny et al., 1988; Hayden and Gallant, 2009; Khayat et al., 2010; Martinez-Trujillo and Treue, 2004; McAdams and Maunsell, 2000; Motter, 1994; Treue and Martinez Trujillo, 1999), objects (Blaser et al., 2000; Houtkamp et al., 2003; Serences et al., 2004), and periods (Coull and Nobre, 1998; Doherty et al., 2005; Ghose and Maunsell, 2002).

Whether all forms of attention employ common neural mechanisms has been debated extensively (Duncan, 1980; Maunsell and Treue, 2006). Several psychophysical studies have argued that spatial attention is unique and that nonspatial forms of attention are inextricably tied to spatial location (Kwak and Egeth,

1992; Nissen and Corkin, 1985). However, other studies argue that spatial and nonspatial forms of attention are qualitatively similar and might be mediated by equivalent mechanisms (Bundesen, 1990; Duncan, 1980; Keren, 1976; Rossi and Paradiso, 1995; von Wright, 1970).

Neurophysiological studies provide evidence supporting both views. Both spatial attention (Assad, 2003; Maunsell and Treue, 2006; Reynolds and Chelazzi, 2004; Yantis and Serences, 2003) and feature attention (Assad, 2003; Hayden and Gallant, 2009; Martinez-Trujillo and Treue, 2004; Maunsell and Treue, 2006; McAdams and Maunsell, 2000; Motter, 1994; Reynolds and Chelazzi, 2004; Treue and Martinez Trujillo, 1999; Yantis and Serences, 2003) modulate the responses of individual sensory neurons: attending to a stimulus or feature that matches a neuron's receptive field location or tuning preference typically increases neuronal responses. The similarity in the way different forms of attention affect individual neurons led to the hypothesis that all forms of attention use a similar neuronal mechanism (Martinez-Trujillo and Treue, 2004; Maunsell and Treue, 2006; Treue and Martinez Trujillo, 1999).

However, the retinotopic organization of visual cortex may allow spatial attention to employ a distinct mechanism because the comodulated neurons are typically located near each other. Spatial attention may be mediated by feedback from pre-motor cells in the frontal and parietal areas involved in eye movement planning (for review, see Astafiev et al., 2003; Bisley and Goldberg, 2010; Craighero et al., 1999; Gitelman et al., 1999; Moore et al., 2003); such feedback may target local groups of neurons.

In contrast, most features are represented by neurons that are dispersed throughout cortex. Attending to these features would require a mechanism that does not rely on topographic organization. One possibility is that attention to such features is only possible through learning and longer-term plasticity (Wolfe et al., 2004), and all forms of attention may require topographic organization. Perhaps because attention to topographically organized features is more natural, most neurophysiological studies have focused on attention to topologically organized features, most notably motion direction in the middle temporal area (Albright, 1984; Martinez-Trujillo and Treue, 2004; Sally et al., 2009).

Over blocks of behavioral trials, the attentional modulation of either behavior or neuronal responses depends largely on the details of the behavioral paradigm chosen by experimenters. However, cognitive states such as attention inevitably fluctuate from trial-to-trial, even within a task condition. We showed recently that the responses of populations of sensory neurons

can be used to detect trial-to-trial fluctuations in spatial attention that are predictive of psychophysical performance (Cohen and Maunsell, 2010). These spontaneous attentional fluctuations can provide hints about the mechanisms mediating feature and spatial attention. For example, if feature attention relies on spatial attention to affect behavior (Kwak and Egeth, 1992; Nissen and Corkin, 1985), then fluctuations in feature attention might either covary with fluctuations in spatial attention or else have little effect on behavior relative to fluctuations in spatial attention. Fluctuations in attention can also be used to determine whether either form of attention acts selectively on local groups of neurons by examining the extent to which fluctuations in feature or spatial attention are coordinated across cortex.

We investigated whether spatial and feature attention employ common or unique mechanisms by analyzing the responses of populations of neurons in visual area V4 in both cerebral hemispheres. We found many qualitative and quantitative similarities between the two types of attention, including their effects on local populations of neurons and the extent to which they could be estimated on individual trials from the responses of a few dozen neurons, suggesting that they employ similar neuronal mechanisms. However, we found that unlike spatial attention, which targets spatially localized groups of neurons in V4, feature attention selectively comodulates neurons located far apart, even in opposite hemispheres. Our results are consistent with the idea that feature and spatial attention are separate processes that rely on similar mechanisms. These results also provide a constraint on a general attentional mechanism: it must be able to modulate the responses of specific and arbitrary subgroups of neurons, even when they are located far apart in cortex.

RESULTS

We trained two rhesus monkeys (*Macaca mulatta*) to perform a change detection task in which we simultaneously manipulated spatial and a form of task or feature attention (Figure 1A). On each trial, two achromatic Gabor stimuli flashed synchronously. At an unsignaled and randomized time, either the orientation or the spatial frequency of one of the stimuli changed. The monkey was rewarded for making an eye movement to the stimulus that changed within 500 ms. We manipulated attention by cueing the monkey in blocks as to which of the two stimuli was more likely to change (left or right: spatial attention) and which stimulus feature would change (orientation or spatial frequency: feature attention; see Experimental Procedures).

We only included data sets in which the monkey completed at least four blocks of each spatial and feature attention condition. Spatial attention alternated on successive blocks and feature attention alternated every four blocks (Figure 1B). We attempted to choose ranges of orientation and spatial frequencies so that the animals' average performance in the two tasks was equivalent (overall performance for the two animals on the orientation task was 64% correct, 8% standard deviation [SD]; 92% correct, 2% SD at the largest change; overall performance on the spatial frequency task was 68% correct, 11% SD; 95% correct, 4% SD at the largest change).

While animals performed this change detection task, we recorded simultaneously from all the extracellular microelectrodes

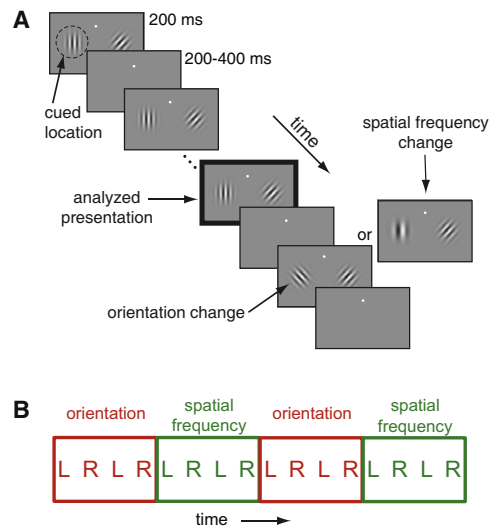


Figure 1. Behavioral Task

(A) Schematic of the orientation and spatial frequency change detection task. Unless otherwise stated, all analyses were performed on responses to the stimulus before the orientation or spatial frequency change (black outlined panel).

(B) Example attention block structure. Spatial attention alternated every block, and feature attention alternated every four blocks. Each data set contained at least four sets of eight blocks (twice as many blocks as depicted here).

in a 6×8 array in V4 in each cerebral hemisphere. The data presented here are from 9 days of recording. We recorded from a total of 68 single units and 588 multiunits. We did not find any significant differences in the effect of attention on single and multiunits (see also Cohen and Maunsell, 2009) and many of the analyses presented here required large simultaneously recorded neuronal populations, so single and multiunits are combined for all analyses.

Task-Related Attention Modulates Single Neurons Similarly to Previously Studied Forms of Feature Attention

The type of task-based feature attention that we used differs from previous studies that manipulated feature attention by changing the visual stimulus outside the neuron's receptive field (Hayden and Gallant, 2009; Martinez-Trujillo and Treue, 2004; Treue and Martinez Trujillo, 1999). We directed the animals to pay attention to either orientation or spatial frequency, rather than one orientation versus another. Also, in our task, there were no visual differences between attention conditions during the period in which we analyzed responses. We focused all analyses on the stimulus presentation immediately before the change, when the stimuli were identical in every trial. The only difference between attention conditions was the location and type of stimulus change the animal was expecting.

We first verified that this type of feature attention affects individual neurons in the same way as other types of feature attention. We quantified the effect of feature attention on each neuron's responses using a standard modulation index that

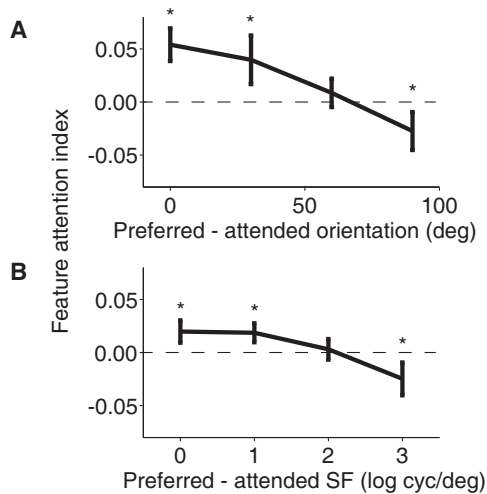


Figure 2. Feature Attention Modulates the Gains of Individual Neurons

(A) Feature attention index as a function of the difference between the neuron's preferred orientation and the orientation of the repeating stimulus. Error bars represent standard error of the mean (SEM). Stars indicate bins for which the average attention index was significantly different than 0 (t test, $p < 0.05$).

(B) Same for spatial frequency. This feature attention index has opposite sign as the index in Figure 2A.

measured the difference between mean responses divided by the sum. We obtained orientation and spatial frequency tuning data by measuring responses to Gabor stimuli with the same size and position as those used in the main task and varying orientation and spatial frequency (see [Experimental Procedures](#)). We selected neurons that showed at least a 2:1 ratio of mean responses to the preferred and orthogonal orientations (147 of 656 neurons; [Figure 2A](#)) or best and worst spatial frequency (314 of 656 neurons; [Figure 2B](#)). We found that neurons whose preferred orientation ([Figure 2A](#), left) or spatial frequency ([Figure 2B](#), left) matched the repeating stimulus before the change showed positive attention indices. This means that, as predicted by the feature-similarity-gain-model ([Martinez-Trujillo and Treue, 2004](#)), attention increases firing rates for neurons whose tuning matches the attended feature. Conversely, we found that feature attention decreased the responses of neurons whose tuning did not match the attended stimulus ([Figure 2A](#) and [2B](#), right). The negative attention indices in the right side of [Figure 2A](#), for example, indicate that attending to a nonpreferred orientation decreases firing rates relative to attending to an average spatial frequency.

Spatial and Feature Attention Affect Local Populations of Cells in Similar Ways

Whereas both feature and spatial attention are known to modulate the gains of individual neurons, the effect of feature attention on the local interactions between neurons is unknown. We showed previously that in addition to increasing the mean responses of individual neurons, spatial attention decreases correlations between neurons in the same hemisphere ([Cohen and Maunsell, 2009](#)). If both forms of attention employ the

same mechanism, feature attention should modulate correlations between nearby neurons as well.

We quantified the extent to which the trial-to-trial fluctuations in the responses of a pair of neurons were correlated using a standard measure of spike count correlation (also called noise correlation). For each pair of simultaneously recorded neurons in the same hemisphere, we calculated the Pearson's correlation coefficient of the spike count responses in each attention condition. As in previous studies ([Cohen and Maunsell, 2009](#); [Mitchell et al., 2009](#)), we found that spatial attention modulates correlations, and modulation of rate and correlation are linked ([Figure 3A](#)). The neuron pairs that showed the largest attentional increases in firing rate also showed the biggest decreases in correlation ([Figure 3A](#), upper right). When a pair of neurons showed very little firing rate modulation due to attention, it also typically showed very little change in correlation. Most of the second through fourth quadrants of this plot are empty because few neurons have their rate of firing strongly reduced by spatial attention.

We found that like spatial attention, feature attention affects both rates and correlations and that the magnitudes of these effects covary. Although spatial attention increases the firing rates of most neurons ([Figure 3A](#)), feature attention can either increase or decrease firing rates ([Figure 2](#)). The presence of both positive and negative rate modulations gives us further dynamic range to test the hypothesis that modulations in firing rate correspond to opposite modulations in correlation. In the plot in [Figure 3B](#), we arbitrarily define positive rate changes as stronger responses when the animal was performing the orientation rather than the spatial frequency change detection task. The plot verifies that, as in spatial attention, pairs of neurons whose firing rates increase with feature attention show decreases in correlation ([Figure 3B](#), top right). Conversely, neurons whose firing rates decreased with feature attention showed increases in correlation ([Figure 3B](#), bottom left).

The relationship between modulation of rate and of correlation was quantitatively similar for the two types of attention ([Figure 3C](#)). The slopes of the best fit lines relating the change in noise correlation for each pair to their mean modulation of firing rate were statistically indistinguishable for feature attention (slope, -0.0036 , 95% confidence interval [CI] -0.0058 to -0.0014 ; 12,162 same-hemisphere pairs with similar modulation; see [Experimental Procedures](#)) and spatial attention (slope -0.0037 , 95% CI -0.0049 to -0.0024 ; 63,656 same-hemisphere pairs with similar modulation). The y-intercepts of the best-fit lines were also indistinguishable from each other and from zero (feature intercept = -0.010 ± 0.025 , 95% CI, spatial intercept = -0.001 ± 0.009 , 95% CI).

In principle, we could have obtained the results in [Figures 3A–3C](#) if the rates and correlations of separate populations of cells were modulated by spatial and feature attention. Instead, we found that most cells were modulated to some extent by both types of attention. [Figure 3D](#) shows how modulations by spatial and feature were distributed among cells. No separate subpopulations are obvious.

Our data suggest that spatial and feature attention affect local populations of cells in similar ways. Both types of attention modulate the firing rates of individual neurons as well as pairwise

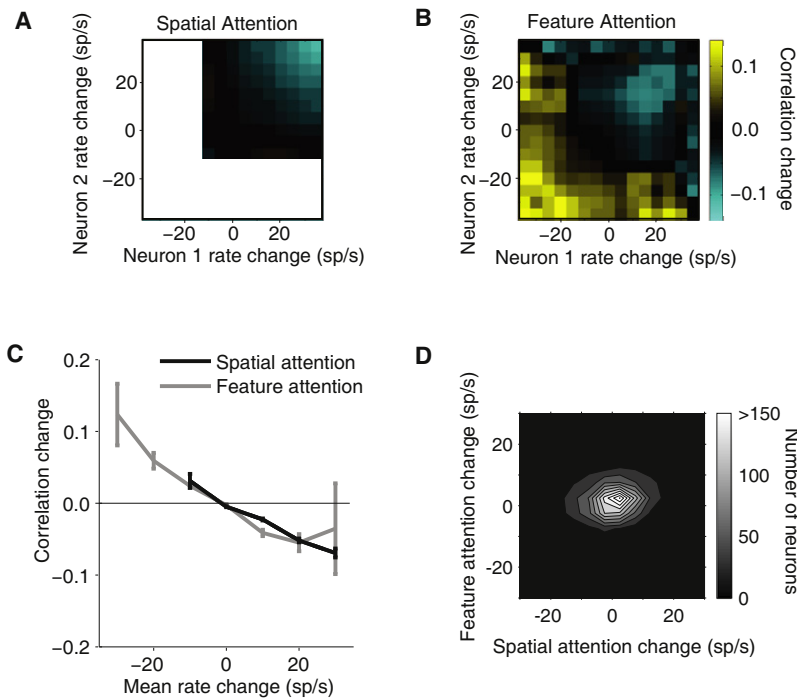


Figure 3. Spatial and Feature Attention Affect Correlations and Firing Rates in Similar Ways

(A) Effects of spatial attention on firing rates and spike count correlations. The x and y axes represent rate changes due to attention for each pair of neurons recorded simultaneously in the same hemisphere (mean response to the stimulus preceding the stimulus change when attention is directed to the contralateral hemifield minus mean rate when attention is directed to the ipsilateral hemifield; 68,846 pairs). Colors represent the change in spike count correlation due to attention (contralateral minus ipsilateral). The second through fourth quadrants of this plot are largely empty because few neurons have their rate of firing strongly reduced by attention. The data are reflected across the diagonal.

(B) Same as (A) for feature attention (16,696 pairs). Positive values indicate higher rates or correlations in the orientation than the spatial frequency change detection task.

(C) Correlation change versus firing rate change for pairs of neurons whose modulation by spatial (black line) or feature attention (gray line) differed by <5 spikes/s. Error bars represent SEM.

(D) Contour plot of rate modulation by feature attention as a function of modulation by spatial attention.

spike count correlations. The tight link (and inverse relationship) between attentional modulation of rates and correlations suggests that both changes may be mediated by a single mechanism that decreases correlations whenever gains are increased.

Single Trial Measures of Feature and Spatial Attention Reliably Predict Performance on Individual Trials

Like all neuronal and behavioral processes, attention varies from moment to moment. Analyzing attentional fluctuations is revealing for three reasons. First, we can determine whether, like spatial attention (Cohen and Maunsell, 2010), fluctuations in feature attention are identifiable from the responses of a few dozen cells and are associated with changes in psychophysical performance. Second, we can address the question of whether feature attention is dissociable from spatial attention by determining whether, for a given spatial attention state, fluctuations in feature attention affect behavior. Finally, fluctuations in attention can reveal the cortical extent of modulation by either form of attention. If distant groups of neurons are comodulated by attention, then the strength of their attentional modulation should be correlated on a trial-to-trial basis.

These analyses require an estimate of the animal's attentional state on a single trial. An instantaneous measure of spatial attention based on the responses of populations of V4 neurons can reliably predict an animal's ability to perform a difficult psychophysical task several hundred milliseconds in the future (Cohen and Maunsell, 2010). We used this measure and an analogous measure of feature attention to predict behavior to examine spatial extents of the two types of attention.

Our task had four attention conditions: each trial belonged to one of two spatial attention conditions (left or right) and one of

two feature attention conditions (orientation or spatial frequency). Using similar methods to those in our previous study (Cohen and Maunsell, 2010), we quantified attention on a single trial as the similarity of the population response to the mean responses in each attention condition. This method is not an ideal decoder to distinguish between correct and incorrect trials based on population responses. Instead, we tested the hypothesis that a single-trial extension of the traditional definition of attention, which compares mean responses in different attention conditions (e.g., Figure 2) could predict behavior. We focused our analyses on trials with a single, difficult orientation change or a single, difficult spatial frequency change for which all trials had valid attentional cues. The average performance on these trials was 34% correct across all data sets (total correct trials divided by total correct plus total missed trials), which is in a range where attention can be the difference between correct and incorrect trials.

We first plotted the population response on each trial in an n -dimensional space in which each of the n simultaneously recorded neurons represented one dimension. If we recorded 83 neurons in the two hemispheres combined, the population response on each trial would be a point in an 83-dimensional space. For ease of visualization, we have plotted these responses for two simultaneously recorded neurons in an example recording session (in a two-dimensional space; Figures 4A and 4C), but the actual analyses used all simultaneously recorded neurons in a high-dimensional space.

We then projected each response onto a putative "spatial attention axis" and a putative "feature attention axis" using a process that is illustrated for the data from an example recording session in Figures 4A–4D. The spatial attention axis for a given trial was the line in the n -dimensional space that

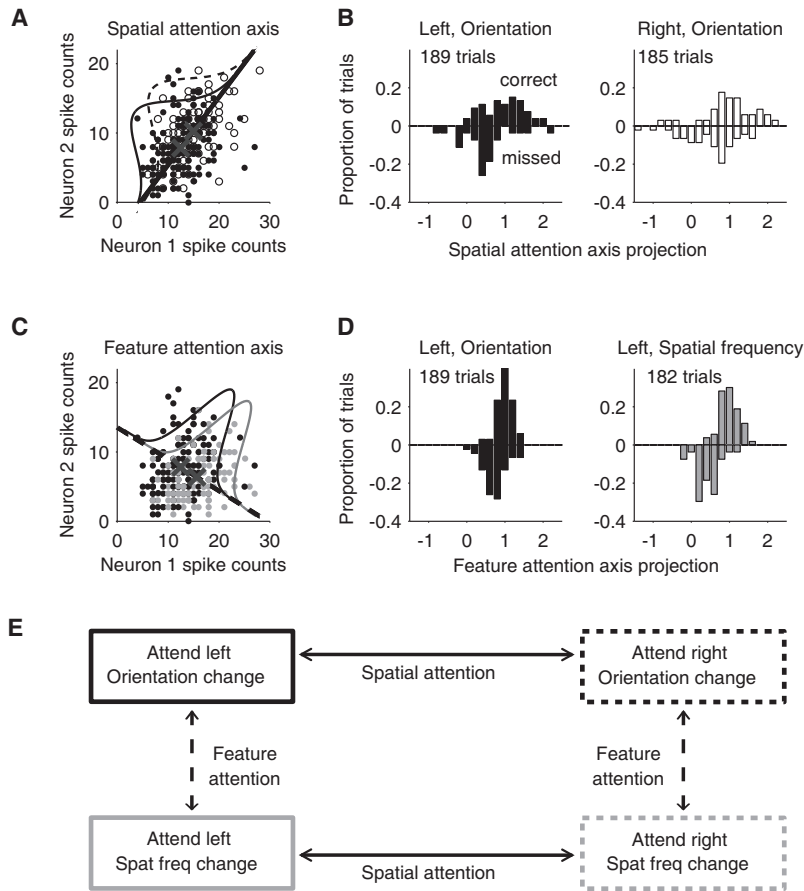


Figure 4. Estimates of Both Feature and Spatial Attention Predict Behavior on Individual Trials

(A) Procedure for calculating allocation of spatial attention on a single orientation change trial. For each trial, the number of spikes fired by n simultaneously recorded neurons during the stimulus before an orientation change in the left hemifield (open points) and right hemifield (filled points) is plotted as a point in an n -dimensional space (a two-neuron example showing unusually large attention effects is plotted here). The spatial attention axis (black line) is the line connecting the center of mass of the n -dimensional cloud of points for correct trials at each attention/change location (X). Each point (including missed trials) is projected onto the axis. The projections are scaled for each data set so that a projection of $+1$ is equal to the mean response before correct detections in the same attention condition as a given trial and -1 is equal to the mean before correct detections in the opposite attention condition.

(B) Frequency diagram of population projections on trials with left (left plot) or right orientation changes (right plot) for the same example day before correct detections (upward bars) and missed changes (downward bars).

(C) Same as (A), for the feature attention axis comparing orientation (black points) and spatial frequency change detections (gray) on attend-left trials. The attend-left, orientation change trials (black points) are the same as in (A).

(D) Same as (B), for feature attention. The attend-left, orientation change trials (black bars) are the same as in (B), but here are projected onto a feature attention axis rather than a spatial attention axis.

(E) Construction of feature and spatial attention axes. Each trial had both a spatial attention condition (attend left or right) and a feature attention condition (orientation or spatial frequency change), and so belonged to one of four conditions. The four conditions were used to construct four attention axes, two of which were relevant for a given trial.

connected the mean responses before correct detections in the two spatial attention conditions that had the same feature attention condition as that trial (Figure 4A, and Figure 4E, horizontal axes). For example, the population responses in trials in the attend-left, orientation change condition (Figure 4A, black points) are projected onto the spatial attention axis connecting mean responses in the attend left and attend right conditions in the orientation change detection task (Figure 4E, top horizontal axis). Similarly, the feature attention axis connected the mean responses before correct detections in the two feature attention conditions that had the same spatial attention condition as the given trial (Figure 4C, dashed line; Figure 4E, vertical axes).

These projections provide two simultaneous measures of attention for each trial: an estimate of spatial attention and an estimate of feature attention. To compare across recording sessions, we normalized the scalar projections onto the two axes for each recording session so that a projection of $+1$ was equal to the mean response before correct detections in the same attention condition as a given trial, and -1 was equal to the mean response in the opposite condition. These projections are plotted in Figures 4B and 4D for the same example recording session as the example neurons in Figures 4A and 4C (although the projections are computed from the responses of all 83 neurons that were

simultaneously recorded during that session). Each trial has a projection on both a spatial and a feature attention axis.

The attention axes were defined based on population responses in only correct trials. Because of the way we normalized the projections, the means of all distributions of projections for correct trials are by definition $+1$. Responses on missed trials provide an independent test of the hypothesis that position on the attention axis predicts behavioral performance. The specific hypothesis is that missed detections are more likely to occur when the projection on the attention axis moves from the mean of the correct attention condition toward the opposite attention condition. The mean of correct trials in the opposite condition was normalized to be -1 , so projections less than $+1$ indicate less attention was allocated to the appropriate location or feature. Consistent with this hypothesis, the means of all of the distributions of missed trials for the example recording session were <1 (Figures 4B and 4D), indicating that behavioral performance correlated with position on both the spatial and feature attention axes. Other methods of defining the axes (e.g., using half of all trials, or half of correct trials) produced qualitatively similar results.

By normalizing the projections for each recording session and attention axis, every trial can be assigned to a point in

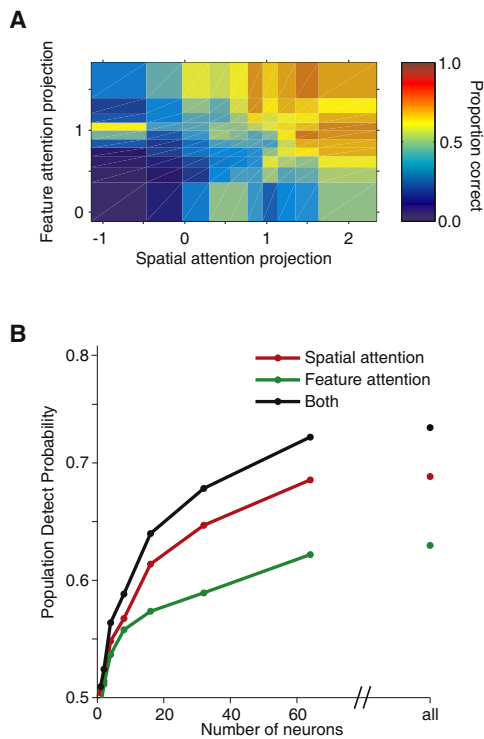


Figure 5. Behavioral Performance Depends on Projections onto Both the Feature and Spatial Attention Axes

(A) The middle 99% of projections were divided into 10 equally sized feature attention bins and 10 equally sized spatial attention bins. The colors represent the proportion correct (number of correct trials divided by the sum of correct and missed trials).

(B) Population DP_{AA} as a function of number of neurons. The rightmost points represent DP_{AA} for all simultaneously recorded cells (mean, 83 single and multiunits).

a two-dimensional plot of spatial attention and feature attention. For example, consider an attend-left, orientation change trial with a spatial attention projection of +1 and a feature attention projection of -1 . These projections mean that the projection of the population response on that trial onto the spatial attention axis connecting the means of correct attend-left and attend-right trials in the orientation change detection was equal to the mean projection for correct attend-left, orientation change trials. The feature attention projection of -1 means that the projection of that same population response onto the axis connecting the mean responses on attend-left orientation change and attend-left spatial frequency trials was equal to the mean projection in the opposite condition (attend-left spatial frequency trials in this example).

Across our recording sessions, behavioral performance correlated strongly with position on the spatial attention axis (Cohen and Maunsell, 2010) and the feature attention axis (Figure 5A). We discarded the outlying 1% of trials on each axis (0.5% of trials with the largest and smallest projections onto each axis; 1.96% of total trials) and assigned the remaining trials to a bin based on position on the spatial attention axis (x axis) and the feature attention axis (y axis) such that 10% of the remaining

data was in each bin. The color of each bin represents the animal's proportion correct for each combination of projections onto the spatial and feature attention axes.

We observed substantial variability along both axes. The mean projections for correct trials were defined as +1. Spatial attention varied from >2 to -1 (that corresponds to the mean of the opposite spatial attention condition) on this scale. Feature attention varied less, from 1.5 to 0. The lower variability along the feature axis was likely caused by the less frequent feature attention block changes (Figure 1B). Also, feature attention cues were always valid whereas changes sometimes occurred at the uncued location, encouraging the animal to direct some attention there.

The trial-to-trial variability in both spatial and feature attention was associated with large changes in behavior. Performance on trials in which the animal's attention was directed strongly toward the correct feature (Figure 5, top row) or correct location (Figure 5A right column) was much better than when the animal's attention was only weakly directed toward the correct feature or location (bottom row and left column, respectively). The average performance for the four bins in the upper right of Figure 5A was 71% correct (95% CI, 63% to 78% correct), whereas the average performance for the four bins in the lower left was 10% correct (95% CI, 6% to 14% correct).

Estimates of Both Feature and Spatial Attention Predict Behavior on Individual Trials

We summarized the relationship between attention axis position and performance by calculating the area under the receiver operating characteristic (ROC) curve for the distributions of positions before correct and missed detections. This measure is comparable detect probability (DP), which has been used to quantify the ability of an ideal observer to predict an animal's behavioral choice based on the responses of single sensory neurons to the changed stimulus (Cohen and Maunsell, 2010; Cook and Maunsell, 2002), called choice probability for a discrimination task (Britten et al., 1996; Parker and Newsome, 1998). Our metric differs in that it is based on population projections onto an attention axis rather than spike counts from single neurons and in that it relies on responses to stimuli before the stimulus change. We refer to our metric as DP_{AA} to emphasize that this calculation is done on projections onto the attention axis (AA) (Cohen and Maunsell, 2010).

As Figure 5A suggests, both feature and spatial attention predict performance, although spatial attention was more predictive. The average DP_{AA} for feature attention was 0.63, and DP_{AA} for spatial attention was 0.68. This measure was significantly greater than 0.5 for both types of attention (t tests; $p < 10^{-3}$).

We assessed the dependence of DP_{AA} on the number of neurons from which the attention axis projections were calculated (Figure 5B). For each recording session, we randomly selected (without replacement) subsets of neurons, calculated projections onto an attention axis constructed for just those neurons, computed the area under the ROC curve comparing the distributions of projections for correct and missed trials, and repeated the process 1000 times. For the combined feature and spatial attention axes, we calculated the percent correct

classifications of the ideal linear discriminator between the two-dimensional distributions of projections for correct and missed trials. DP_{AA} increases with population size, and appears to approach asymptote at population sizes only slightly larger than our mean of 83 neurons.

Fluctuations in Attention, Rather Than in Global Factors, Predict Behavior

We used this metric to test the possibility that some of the variability along the attention axis arose from variability in global factors such as arousal or alertness rather than variability in attention. This possibility seems unlikely, because both attention axes should be orthogonal to global axes. About half the neurons increase their rates and half decrease their rates in each attention condition. For spatial attention, neurons with receptive fields in the left hemifield tend to have higher firing rates in the attend-left than the attend-right condition, and the opposite is true for neurons whose receptive fields are in the right hemifield. For feature attention, about half of the neurons in each hemisphere respond more strongly in the orientation change than the spatial frequency change detection task.

In contrast, global factors should comodulate all neurons. To directly test the possibility that global factors can predict behavior, we computed projections onto a response axis (from the origin to the mean response to the repeated stimulus). The response axis did not predict behavior well; average DP was slightly above chance for the left hemisphere ($DP = 0.53$; *t* test, $p < 0.05$) and indistinguishable from chance for the right hemisphere ($DP = 0.51$; *t* test, $p = 0.14$). In each hemisphere, DP for the response axis was significantly different than DP_{AA} for either attention axis (paired *t* tests, $p < 0.01$). These results suggest that fluctuations in global factors do not account for the ability of the feature and spatial attention axes to predict behavior.

Feature, but Not Spatial, Attention Is Coordinated across Hemifields

The ability to estimate attention on individual trials can also provide insight into the cortical extent of modulation by spatial and feature attention. We showed that fluctuations in the amount of attention allocated to two stimuli in opposite hemifields is uncorrelated, suggesting that spatial attention is mediated by retinotopically local processes (Cohen and Maunsell, 2010). We replicated this result for the current data set by defining spatial attention axes separately for neurons recorded from the two arrays (corresponding to neurons whose receptive fields are in opposite hemifields). The projections onto each axis were thus independent estimates of attention allocated to each stimulus. We calculated the correlation between the projections onto the two axes within each attention condition (Figure 4E).

The correlation between projections on the spatial attention axes for the two cerebral hemispheres was indistinguishable from 0 (Figure 6A, black bar; *t* test, $p = 0.24$). This lack of correlation was not a result of insufficient statistical power: when we randomly divided the neurons recorded within a hemisphere into two equal-sized groups, we easily detected a positive correlation between projections onto spatial attention axes calculated from each subgroup (Figure 6B, black bar; $p < 10^{-10}$). Our data indicate that fluctuations in the amount of spatial attention allo-




Figure 6. Fluctuations in Feature, but Not Spatial Attention Are Coordinated across the Two Hemispheres

(A) Mean correlation coefficient between population projections onto attention axes constructed using simultaneously recorded neurons in the two hemispheres. The mean correlation coefficient was statistically >0 for feature attention (*t* test, $p < 10^{-6}$) and indistinguishable from zero for spatial attention ($p = 0.24$) or the correlation between spatial and feature attention ($p = 0.09$). (B) Same for randomly chosen subsets of neurons within a hemisphere. The mean correlation coefficients were statistically >0 for both spatial and feature attention ($p < 10^{-10}$) and indistinguishable from zero for the correlation between spatial and feature attention ($p = 0.16$).

cated to the two stimuli arise from fluctuations in groups of neurons within a hemisphere, rather than because the animal attends to the wrong stimulus.

The cortical extent of feature attention is qualitatively different. As before, we constructed a separate feature attention axis for neurons in each hemisphere and calculated the correlation coefficient between projections on the two axes. Our statistical power for detecting correlations along the feature attention axes was similar for feature and spatial attention (Figure 6B; *t* test for feature attention, $p < 10^{-10}$). However, in contrast to spatial attention, we found that projections on the two feature attention axes were positively correlated across hemispheres (Figure 6A, gray bar; $p < 10^{-6}$).

We did not find evidence that fluctuations in feature attention are linked to fluctuations in spatial attention. The correlation between fluctuations in spatial and feature attention is indistinguishable from 0 both across hemispheres (Figure 6A, white bar; $p = 0.09$) and within a hemisphere (Figure 6B, white bar; $p = 0.16$).

Fluctuations in the Responses of Attentionally Modulated Neurons in Opposite Hemispheres Are Correlated

The positive correlation between the positions on the feature attention axes in the two hemispheres is in apparent conflict with the finding that spike count correlations between pairs of neurons in opposite hemispheres are weak (Cohen and Maunsell, 2009). As in our previous study, average spike count correlations between opposite hemispheres in this data set were small (mean = 0.017, SD = 0.09). The positive correlation between the fluctuations in feature attention to the two stimuli, however, can only come about from confluctuations between neurons in opposite hemispheres. These observations suggest that the neurons whose responses contribute most to the feature attention axes are positively correlated while those that contribute most to the spatial attention axes are on average uncorrelated.




Figure 7. Fluctuations in the Responses of Comodulated Neurons in Opposite Hemispheres Are Correlated

Spike count correlation is plotted as a function of attention modulation for pairs of neurons in opposite hemispheres whose absolute value modulation by spatial (black lines) or feature attention (gray lines) was within 5 spikes/s (see text). Solid lines represent pairs with same-sign modulation, and dashed lines represent pairs with opposite sign modulation. The gray numbers represent the number of pairs that contribute to each feature attention bin, and the black numbers represent the pairs that contribute to each spatial attention bin.

The attention axis effectively weights neurons by the magnitude of the attentional modulation of their mean rates. To assess correlations among the neurons that contribute most to each attention axis, we sorted neurons by their mean attentional modulation and selected pairs from opposite hemispheres that had similar mean spatial or feature attention modulation (see [Experimental Procedures](#)). Consistent with our prior observation that on average, pairwise correlations between the hemispheres are small, we observed small correlations between pairs with attention modulation <5 sp/s (which were the majority of pairs; [Figure 7](#)).

In contrast, correlations between pairs of strongly modulated cells depended greatly on whether attention modulated their mean responses in the same or opposite directions. In general, spatial attention modulates opposite hemisphere pairs with opposite sign: attending to the left tends to increase the rates of neurons in the right hemisphere and decrease the responses of neurons in the left hemisphere. Therefore, the modulation of the majority of pairs of neurons in opposite hemispheres will have opposite signs. The fluctuations in the responses of these pairs with opposite-sign spatial attention modulation have near zero correlation, with a slight trend toward lower correlations for pairs with larger modulations ([Figure 7](#), black dashed line). The minority of pairs with same sign modulation (which by definition include one neuron whose response was lower on attended than unattended trials) showed a different pattern of activity. There were no pairs with strong same sign modulation because this would mean that one neuron had a strong “wrong

way” modulation (much lower firing rate when attention was directed to the stimulus in its receptive field than when attention was directed to the opposite hemifield). Nevertheless, pairs with a mean same-sign modulation of 5 or 10 spikes/s had higher spike count correlations than pairs whose modulation was the same magnitude but of opposite sign (black solid line).

The relationship between attentional modulation and correlation is nearly identical for feature attention, but the distributions of pairs with same- or opposite-sign modulation are different. Unlike spatial attention, feature attention increases and decreases neurons in a given hemisphere with about equal probability. Therefore, approximately half of opposite hemisphere pairs show same sign feature attention modulation (i.e., both have higher firing rates during the orientation than the spatial frequency task, or vice versa) and half have opposite sign modulation. As in spatial attention, pairs with opposite sign feature attention modulation have weak correlations ([Figure 7](#), gray dashed line). In contrast, pairs with strong same-sign modulation have strongly positive correlations (gray solid line). These results suggest that neurons that are comodulated by attention share a common input, even when they are in opposite hemispheres.

This observation also explains the differences in the extent to which fluctuations in feature and spatial attention are coordinated across hemispheres ([Figure 6A](#)). Because the attention axis runs through the difference between mean responses in two attention conditions, neurons that are strongly modulated by attention dominate projections onto the axis. Nearly all pairs of neurons in opposite hemispheres that are strongly modulated by spatial attention have opposite-sign modulation ([Figure 7](#)). The fluctuations in the responses of these neurons are nearly uncorrelated, so projections onto the two attention axes are uncorrelated as well ([Figure 6A](#)). In contrast, approximately half of the opposite hemisphere pairs that are strongly modulated by feature attention have same-sign modulation, so the attention axes are dominated at least in part by pairs with positive correlations.

DISCUSSION

We simultaneously manipulated feature and spatial attention to assess their effects on local and spatially disparate populations of neurons. The observation that the two forms of attention vary independently ([Figure 6](#)) allowed us to assess their effects on V4 neurons separately but on the same behavioral trials. Using this task, we replicated the single neuron results of previous studies that manipulated each type of attention separately ([Cohen and Maunsell, 2009; Maunsell and Treue, 2006; Treue and Martinez Trujillo, 1999](#)) and the effects of spatial attention on correlations between nearby neurons ([Cohen and Maunsell, 2009; Mitchell et al., 2009](#)), suggesting that simultaneously manipulating feature and spatial attention employs the same mechanisms as manipulating each separately.

Analyzing the effect of attention on populations of neurons provides several new means of comparing spatial and feature of attention. Here, we review the implications of these data for the hypothesis that the two forms of attention are mediated by a common mechanism and discuss the potential for using population data for understanding the neural circuitry underlying other sensory, motor, and cognitive processes.

Similar but Separate Processes

We found many similarities in the way that spatial and feature attention modulated local populations of neurons and affected behavior, supporting the hypothesis that the two types of attention are mediated by similar mechanisms. Both types of attention affect correlations between pairs of nearby neurons as well as the firing rates of individual neurons (Figure 3). The striking quantitative similarity in the relationship between correlation and rate changes between the two forms of attention (Figure 3C) suggests that a single process modulates both correlation and rate.

Comparing the effects of spontaneous fluctuations in the two forms of attention on behavior allowed us to look beyond interactions between feature and spatial attention that are imposed by the structure of the task. We showed that fluctuations in both forms of attention are responsible for large changes in behavioral performance (Figure 5A). The two types of attention vary independently (Figures 6A and 6B, white bars), and fluctuations in feature attention occur and modulate behavior even when spatial attention is constant (Figure 5). Our results indicate that feature and spatial attention are separable processes, each with the ability to affect psychophysical performance.

The primary difference between spatial and feature attention in our data set is that fluctuations in feature attention are coordinated across hemispheres (Figure 6) and that the responses of pairs that show strong feature attention effects are comodulated on a trial-to-trial basis (Figure 7), whereas spatial attention is independent by both measures. These results are consistent with the idea that spatial attention acts on local groups of neurons, and that the amount of attention allocated to locations in opposite hemifields is independent. In contrast, attention to features appears to be coordinated across the visual field, suggesting that feature attention selectively comodulates neurons located far apart, even in opposite hemispheres.

The idea that spatial and feature attention operate on different spatial scales is supported by psychophysical evidence. A subject's ability to spatially attend to an object in one hemifield is unaffected by attention to objects in the other hemifield (Alvarez and Cavanagh, 2005). Conversely, feature attention can affect visual processing independent of stimulus location (Liu and Mance, 2011; Saenz et al., 2002, 2003).

Possible Attentional Mechanisms

To be consistent with our data, a unified attention mechanism must operate on a more local group of neurons for spatial attention than feature attention. The independence of spatial attention across hemispheres is consistent with the premotor theory of spatial attention. This theory postulates that spatial attention is mediated by feedback from pre-oculomotor neurons (Astafiev et al., 2003; Bisley and Goldberg, 2010; Craighero et al., 1999; Gitelman et al., 1999; Moore et al., 2003), which may target local populations of cells. This theory is supported by evidence showing that microstimulation of areas involved in eye movement planning mimics many of the behavioral and neuronal effects of spatial attention (Cavanaugh et al., 2006; Cavanaugh and Wurtz, 2004; Cutrell and Marrocco, 2002; Herrington and Assad, 2009; Herrington et al., 2009; Moore and Armstrong,

2003; Moore et al., 2003; Moore and Fallah, 2001; Muller et al., 2005).

An analogous mechanism for attention to nontopographically organized features would require flexible feedback from neurons that encode the attended feature, which could potentially come from frontal or late visual areas (such as inferotemporal cortex) that flexibly encode many attributes of visual scenes (for discussion see Maunsell and Treue, 2006). Plasticity may also play a role: feedback connections from frontal areas to the relevant subsets of visual neurons could be strengthened during the training process, consistent with the finding that the ability to attend to complicated patterns and features improves with practice (Wolfe, 1998).

The tight and inverse relationship between attentional modulation of rates and correlations suggests that attention modulates the strength or activity of a common input that reduces the gains of the responses of V4 neurons. A rate increase combined with a correlation decrease is consistent with a decrease in an effectively inhibitory common input. A background input whose role is to reduce the gains of single neurons (Chance et al., 2002) could fill this function. Such inputs could in principle be responsible for the normalization of sensory responses, which may be linked to attention (Lee and Maunsell, 2009; Reynolds and Heeger, 2009; Boynton, 2009). An analogous mechanism for feature attention would require that neurons with similar tuning for the attended feature share a common input that can be selectively modulated by attention. Further work will be needed to determine whether such inputs exist.

Gains, Correlations, and Coding

The precise relationship between gain changes and correlation changes (Figure 3) along with the observations that correlations depend on sensory stimuli (Aertsen et al., 1989; Ahissar et al., 1992; Espinosa and Gerstein, 1988; Kohn and Smith, 2005), learning (Ahissar et al., 1992; Gutnisky and Dragoi, 2008; Komiyama et al., 2010), or other cognitive factors (Cohen and Maunsell, 2009; Cohen and Newsome, 2008; Mitchell et al., 2009; Poulet and Petersen, 2008; Vaadia et al., 1995) support the idea that correlation changes are an important aspect of population coding in cortex. It has long been recognized that correlations affect the amount of sensory information encoded in a population of neurons (Abbott and Dayan, 1999; Averbeck et al., 2006; Shadlen et al., 1996; Zohary et al., 1994).

We showed previously that the reduction in correlations from spatial attention could account for most of the improvement in the amount of sensory information encoded in V4 (Cohen and Maunsell, 2009; Mitchell et al., 2009). Here, we showed that for neurons whose tuning matched the attended feature, feature attention also decreases correlations (Figure 3). Furthermore, as predicted by the feature-similarity-gain model (Martinez-Trujillo and Treue, 2004; Treue and Martinez Trujillo, 1999), feature attention decreases the gains of neurons whose tuning is opposite the attended feature and also increases correlations between these down-modulated cells.

The higher correlations among neurons that are not tuned for the attended location or feature may be a hallmark of neuronal populations that are not well-driven or engaged in a task. A recent study found that trial-to-trial variability in individual neurons in

seven cortical areas is higher when the cells are not well-driven, even after correcting for the expected effects of a lower rate of firing (Churchland et al., 2010). This effect may be a signature of a network in which stimulus drive suppresses correlated ongoing activity (Rajan et al., 2010). At low frequencies, both spike-field coherence and cortical oscillations in the local field potential are higher for populations encoding unattended stimuli (Fries et al., 2001; Gregoriou et al., 2009; Womelsdorf et al., 2007). In humans, withdrawing attention increases low frequency oscillations in MEG signals (Siegel et al., 2008), and functional connectivity (and therefore variability) is often higher during the spontaneous “resting state” than when neural populations are well-driven (for review see van den Heuvel and Hulshoff Pol, 2010).

In contrast, attention increases spike-field coherence and oscillations at high frequencies (Fries et al., 2001; Gregoriou et al., 2009; Womelsdorf and Fries, 2007). These increases have been hypothesized to improve communication between sensory neurons and downstream cells by improving the probability that synchronous spikes will drive a post-synaptic cell above threshold (for review see Womelsdorf and Fries, 2007); but see also (Ray and Maunsell, 2010). Attentional increases in high frequency correlations are not inconsistent with the reductions in low frequency correlations we and others have reported. In principle, the two could work in concert to remove correlations on long timescales while improving neural communication on short timescales.

Cognitive States Are Inevitably Variable

The observation that even in a controlled experimental setting, both spatial and feature attention vary substantially (Figure 5) suggests that all aspects of a subject’s internal state vary from moment to moment and that it is impossible to measure any particular cognitive factor in isolation. The spatial and feature attention axes we defined, which measure differences in the amount of attention allocated to two particular locations and two seemingly nonopposed features, are by no means the only aspects of attention that could vary. The animal may allocate attention to locations other than these two stimuli (e.g., the fixation point or the door to the room) and to features other than orientation or spatial frequency, or other sensory modalities. Other forms of attention, such as those to task timing, and other cognitive processes such as arousal or motivation likely vary and affect behavior as well. The observation that the two attention axes we measured predicted behavior so well indicates that these were important for performance in this task. Further work will be needed to determine the effects of other cognitive processes on sensory neurons and behavior, and the extent to which the influence of each is dependent on the specifics of the task or behavioral context.

Using the Responses of Neuronal Populations to Study Cognitive Processes

In addition to addressing the question of the similarity of feature and spatial attention, our results show that analyzing the relationship between the responses of populations of neurons and behavior can provide new insight into the mechanisms underlying cognitive processes. Simultaneous recordings from popu-

lations of neurons are becoming easier and more popular, but so far, these larger data sets have been used primarily to increase statistical power or to examine correlations between pairs of neurons. We used the responses of all of the neurons we recorded simultaneously to estimate the amount of feature and spatial attention allocated to each stimulus on each trial. These estimates predict behavior on individual trials and are informative about the neuronal mechanisms underlying attention.

Capitalizing on natural fluctuations in cognitive states within a task condition can provide insight about the way cognitive processes affect behavior and about the neuronal mechanisms underlying these processes that are not accessible using other measures. In the current study, we used these methods to investigate interactions between the behavioral effects of feature and spatial attention as well as the cortical extent of modulation by each type of attention. This information is not available in average responses across task conditions: the structure of the task affects the way that the two types of attention modulate behavior and can also impose blockwise correlations between the amount of attention allocated different locations and features. For example, because exactly one stimulus changed per trial and the identity of the stimulus most likely to change alternated between blocks of trials, our task (and many other behavioral tasks) imposes a blockwise anticorrelation in the average amount of spatial attention allocated to the two stimuli. In contrast, the attention axis method revealed that the amount of attention allocated to each stimulus is in fact independent. Furthermore, looking at the effects of feature and spatial attention on individual trials resolved the question of whether feature and spatial attention are separable by revealing that feature attention modulates behavior even when spatial attention is constant and that either form of attention can dominate behavior.

Finally, looking at the relationship between population activity and behavior provides the statistical power to associate the responses of particular groups of neurons with behavior. Correlations between fluctuations in the responses of individual neurons and perceptual decisions (“choice probability”) (for review see Nienborg and Cumming, 2010; Parker and Newsome, 1998) are measurable, but they are often too weak and variable to be useful for distinguishing the contributions of different neurons to a given behavior. In the future, population-based measures may provide a useful way of assessing the contribution of different neuronal cell types or neurons in different cortical areas or circuits to particular behaviors.

EXPERIMENTAL PROCEDURES

Subjects and Electrophysiological Recordings

Our subjects were the same two adult male rhesus monkeys (*Macaca mulatta*, 9 and 12 kg) used in our previous experiments (Cohen and Maunsell, 2009, 2010). All procedures were approved by the Institutional Animal Care and Use Committee of Harvard Medical School. Before training, each animal was implanted with a head post and a scleral search coil for monitoring eye movements. After the animal learned the behavioral task (3–4 months) we implanted a 6 × 8 array of microelectrodes (Blackrock Microsystems) in V4 in each cerebral hemisphere. Each electrode was 1 mm long and the distance between the centers of adjacent electrodes was 400 μm. The two arrays were connected to a percutaneous connector that allowed electrophysiological recordings.

We implanted the arrays between the lunare and superior temporal sulci, which were visible during surgery. The centers of the spatial receptive fields for both monkeys were in the lower hemifield (eccentricities Monkey 1: 3°–5° left hemifield, 5°–8° hemifield; Monkey 2: 10°–15° left hemifield, 15°–30° right hemifield). Monkey 2 underwent an unplanned explantation of both arrays before recordings began, so we implanted new arrays several millimeters dorsal to the sites of the original implants. Consequently, Monkey 2 had more eccentric and more dispersed receptive fields than Monkey 1. The receptive field distributions were the only physiological results that were distinguishable between the two monkeys.

The data presented here are from 9 days of recording in which we obtained sufficient data from both tasks (see below; four data sets from Monkey 1 and five from Monkey 2). We recorded a total of 68 single units and 588 sorted multiunits. All spike sorting was done manually following the experiment using Plexon's Offline Sorter.

Tasks and Behavior

We trained both monkeys to perform a change detection task in which we manipulated spatial and feature attention (Figure 1A). A trial began when the monkey fixated a central spot of light, and he was required to maintain fixation within a 1.5° square window. Two achromatic Gabor stimuli whose size, location, orientation, and spatial frequency were optimized for a single neuron recorded in each hemisphere flashed synchronously on (for 200 ms) and off (for a randomized 200–400 ms interval picked from a uniform distribution). At an unsignaled and randomized time picked from an exponential distribution (minimum, 1000 ms; mean, 3000 ms; maximum, 5000 ms), either the orientation or the spatial frequency of one of the stimuli changed. The monkey was rewarded for making an eye movement between 100 and 500 ms following the change to the stimulus that changed. If no stimulus change occurred within 5000 ms, the monkey was rewarded simply for maintaining fixation. These catch trials were not included for analysis.

We manipulated spatial and feature attention by cueing the monkey in blocks as to which of the two stimuli was more likely to change (spatial attention) and which feature would change (feature attention). Before each block of trials, the monkey performed 10 instruction trials in which only a single stimulus appeared in the location and with the type of change (orientation or spatial frequency) that would change most often in the upcoming block of trials. Instruction trials were not considered in the analysis. Of the 125 trials per block, 25 randomly interleaved trials contained changes in the uncued stimulus. Only one stimulus change occurred in each trial, and the monkey was rewarded for correctly detecting a change in either stimulus, regardless of the cued location. We only included data sets for which the monkey completed at least four blocks of each spatial and feature attention condition and achieved at least 90% correct detections of the easiest orientation and spatial frequency changes.

Importantly, the stimuli preceding the orientation or spatial frequency change were the same on every trial throughout an entire day of data, regardless of the attention condition or eventual stimulus change. We were interested in the effects of attention independent of sensory responses. We therefore focused our analyses on the stimulus presentation immediately before the change because the stimuli were the same at this point on every trial and because the monkey's attentional state at this time was most likely to affect his ability to successfully detect the upcoming change. All of the primary analyses are based on spike count responses calculated from the period between 60 and 260 ms after stimulus onset.

We obtained tuning data for all of the neurons we recorded by measuring responses to a variety of Gabor stimuli either before or after the primary experiments each day. The monkeys performed a single stimulus version of the usual orientation change detection task on a stimulus in the upper visual field (far outside the receptive fields of the neurons under study). At the same time, we synchronously flashed an additional Gabor stimulus in the lower visual field in each hemifield for 100 ms each. The test Gabors had the same size and location as the Gabors in the main attention task that day. We varied either the orientation of the test Gabors while keeping the spatial frequency the same as in the orientation change detection task or the spatial frequency while keeping the orientation the same as in the spatial frequency change detection task. We constructed multidimensional tuning curves using spike count

responses during the period from 60 to 160 ms after stimulus onset. To minimize effects of adaptation, we only analyzed responses to stimuli that occurred after the first stimulus and before the changed stimulus in the orientation change detection task.

To obtain the statistical power to make quantitative comparisons between the effects of the two types of attention, the spatial attention data presented in Figure 3 include an additional 41 data sets for which we only obtained data from the orientation change detection task (50 data sets total). Every aspect of the task was identical to the orientation change detection task used in the nine data sets considered here, except that there were no interleaved blocks of the spatial frequency change detection task. These additional data sets have been described elsewhere (Cohen and Maunsell, 2009, 2010).

Attentional Modulation of Rates and Correlations

To quantify attentional modulation of the rates of individual neurons, we either took the difference between the mean responses to the stimulus preceding correct detections in the two attention conditions (Figure 3 and Figure 7) or computed an attention index by normalizing this difference by the sum of the mean responses in the two conditions (Figure 2). By convention, we expressed spatial attention modulation for each neuron as the mean response when attention was cued toward the stimulus in the contralateral hemifield minus the mean during the ipsilateral hemifield condition. We chose to express feature attention as the mean response during the orientation change detection task minus the mean response during the spatial frequency change detection task. We defined pairs of neurons with similar attentional modulation (Figure 3C and Figure 7) as those whose attentional modulation differed by <5 spikes/s (that corresponds to one spike in our 200 ms response window).

We computed spike count correlations as the Pearson's correlation coefficient between spike count responses to the stimulus preceding the changed stimulus on correct trials within an attention condition. The sign of changes in correlation (Figure 3) followed the same conventions as changes in mean firing rate.

ACKNOWLEDGMENTS

We are grateful to Mark Histed, Adam Kohn, Amy Ni, and Douglas Ruff for helpful discussions and comments on an earlier version of the manuscript. This work was supported by NIH grants K99EY020844-01 (M.R.C.) and R01EY005911 (J.H.R.M.) and the Howard Hughes Medical Institute.

Accepted: April 8, 2011

Published: June 22, 2011

REFERENCES

- Abbott, L.F., and Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Comput.* 11, 91–101.
- Aertsen, A.M., Gerstein, G.L., Habib, M.K., and Palm, G. (1989). Dynamics of neuronal firing correlation: modulation of "effective connectivity". *J. Neurophysiol.* 61, 900–917.
- Ahissar, E., Vaadia, E., Ahissar, M., Bergman, H., Arieli, A., and Abeles, M. (1992). Dependence of cortical plasticity on correlated activity of single neurons and on behavioral context. *Science* 257, 1412–1415.
- Albright, T.D. (1984). Direction and orientation selectivity of neurons in visual area MT of the macaque. *J. Neurophysiol.* 52, 1106–1130.
- Alvarez, G.A., and Cavanagh, P. (2005). Independent resources for attentional tracking in the left and right visual hemifields. *Psychol. Sci.* 16, 637–643.
- Assad, J.A. (2003). Neural coding of behavioral relevance in parietal cortex. *Curr. Opin. Neurobiol.* 13, 194–197.
- Astafiev, S.V., Shulman, G.L., Stanley, C.M., Snyder, A.Z., Van Essen, D.C., and Corbetta, M. (2003). Functional organization of human intraparietal and frontal cortex for attending, looking, and pointing. *J. Neurosci.* 23, 4689–4699.
- Averbeck, B.B., Latham, P.E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* 7, 358–366.

- Bisley, J.W., and Goldberg, M.E. (2010). Attention, intention, and priority in the parietal lobe. *Annu. Rev. Neurosci.* 33, 1–21.
- Blaser, E., Pylyshyn, Z.W., and Holcombe, A.O. (2000). Tracking an object through feature space. *Nature* 408, 196–199.
- Boynton, G.M. (2009). A framework for describing the effects of attention on visual responses. *Vision Res.* 49, 1129–1143.
- Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis. Neurosci.* 13, 87–100.
- Bundesen, C. (1990). A theory of visual attention. *Psychol. Rev.* 97, 523–547.
- Cavanaugh, J., Alvarez, B.D., and Wurtz, R.H. (2006). Enhanced performance with brain stimulation: attentional shift or visual cue? *J. Neurosci.* 26, 11347–11358.
- Cavanaugh, J., and Wurtz, R.H. (2004). Subcortical modulation of attention counters change blindness. *J. Neurosci.* 24, 11236–11243.
- Chance, F.S., Abbott, L.F., and Reyes, A.D. (2002). Gain modulation from background synaptic input. *Neuron* 35, 773–782.
- Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13, 369–378.
- Cohen, M.R., and Maunsell, J.H.R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* 12, 1594–1600.
- Cohen, M.R., and Maunsell, J.H.R. (2010). A neuronal population measure of attention predicts behavioral performance on individual trials. *J. Neurosci.* 30, 15241–15253.
- Cohen, M.R., and Newsome, W.T. (2008). Context-dependent changes in functional circuitry in visual area MT. *Neuron* 60, 162–173.
- Cook, E.P., and Maunsell, J.H.R. (2002). Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nat. Neurosci.* 5, 985–994.
- Coull, J.T., and Nobre, A.C. (1998). Where and when to pay attention: the neural systems for directing attention to spatial locations and to time intervals as revealed by both PET and fMRI. *J. Neurosci.* 18, 7426–7435.
- Craighero, L., Fadiga, L., Rizzolatti, G., and Umiltà, C. (1999). Action for perception: a motor-visual attentional effect. *J. Exp. Psychol. Hum. Percept. Perform.* 25, 1673–1692.
- Cutrell, E.B., and Marrocco, R.T. (2002). Electrical microstimulation of primate posterior parietal cortex initiates orienting and alerting components of covert attention. *Exp. Brain Res.* 144, 103–113.
- Doherty, J.R., Rao, A., Mesulam, M.M., and Nobre, A.C. (2005). Synergistic effect of combined temporal and spatial expectations on visual attention. *J. Neurosci.* 25, 8259–8266.
- Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychol. Rev.* 87, 272–300.
- Espinosa, I.E., and Gerstein, G.L. (1988). Cortical auditory neuron interactions during presentation of 3-tone sequences: effective connectivity. *Brain Res.* 450, 39–50.
- Fries, P., Reynolds, J.H., Rorie, A.E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291, 1560–1563.
- Ghose, G.M., and Maunsell, J.H.R. (2002). Attentional modulation in visual cortex depends on task timing. *Nature* 419, 616–620.
- Gitelman, D.R., Nobre, A.C., Parrish, T.B., LaBar, K.S., Kim, Y.H., Meyer, J.R., and Mesulam, M. (1999). A large-scale distributed network for covert spatial attention: further anatomical delineation based on stringent behavioural and cognitive controls. *Brain* 122, 1093–1106.
- Gregoriou, G.G., Gotts, S.J., Zhou, H., and Desimone, R. (2009). Long-range neural coupling through synchronization with attention. *Prog. Brain Res.* 176, 35–45.
- Gutnisky, D.A., and Dragoi, V. (2008). Adaptive coding of visual information in neural populations. *Nature* 452, 220–224.
- Haenny, P.E., Maunsell, J.H.R., and Schiller, P.H. (1988). State dependent activity in monkey visual cortex. II. Retinal and extraretinal factors in V4. *Exp. Brain Res.* 69, 245–259.
- Hayden, B.Y., and Gallant, J.L. (2009). Combined effects of spatial and feature-based attention on responses of V4 neurons. *Vision Res.* 49, 1182–1187.
- Herrington, T.M., and Assad, J.A. (2009). Neural activity in the middle temporal area and lateral intraparietal area during endogenously cued shifts of attention. *J. Neurosci.* 29, 14160–14176.
- Herrington, T.M., Masse, N.Y., Hachmeh, K.J., Smith, J.E., Assad, J.A., and Cook, E.P. (2009). The effect of microsaccades on the correlation between neural activity and behavior in middle temporal, ventral intraparietal, and lateral intraparietal areas. *J. Neurosci.* 29, 5793–5805.
- Houtkamp, R., Spekreijse, H., and Roelfsema, P.R. (2003). A gradual spread of attention during mental curve tracing. *Percept. Psychophys.* 65, 1136–1144.
- Keren, G. (1976). Some considerations of two alleged kinds of selective attention. *J. Exp. Psychol. Gen.* 105, 349–374.
- Khayat, P.S., Niebergall, R., and Martinez-Trujillo, J.C. (2010). Attention differentially modulates similar neuronal responses evoked by varying contrast and direction stimuli in area MT. *J. Neurosci.* 30, 2188–2197.
- Kohn, A., and Smith, M.A. (2005). Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* 25, 3661–3673.
- Komiyama, T., Sato, T.R., O'Connor, D.H., Zhang, Y.X., Huber, D., Hooks, B.M., Gabbito, M., and Svoboda, K. (2010). Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. *Nature* 464, 1182–1186.
- Kwak, H.W., and Egeth, H. (1992). Consequences of allocating attention to locations and to other attributes. *Percept. Psychophys.* 51, 455–464.
- Lee, J., and Maunsell, J.H.R. (2009). A normalization model of attentional modulation of single unit responses. *PLoS ONE* 4, e4651.
- Liu, T., and Mance, I. (2011). Constant spread of feature-based attention across the visual field. *Vision Res.* 51, 26–33.
- Martinez-Trujillo, J.C., and Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr. Biol.* 14, 744–751.
- Maunsell, J.H.R., and Treue, S. (2006). Feature-based attention in visual cortex. *Trends Neurosci.* 29, 317–322.
- McAdams, C.J., and Maunsell, J.H.R. (2000). Attention to both space and feature modulates neuronal responses in macaque area V4. *J. Neurophysiol.* 83, 1751–1755.
- Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2009). Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63, 879–888.
- Moore, T., and Fallah, M. (2001). Control of eye movements and spatial attention. *Proc. Natl. Acad. Sci. USA* 98, 1273–1276.
- Moore, T., and Armstrong, K.M. (2003). Selective gating of visual signals by microstimulation of frontal cortex. *Nature* 421, 370–373.
- Moore, T., Armstrong, K.M., and Fallah, M. (2003). Visuomotor origins of covert spatial attention. *Neuron* 40, 671–683.
- Motter, B.C. (1994). Neural correlates of feature selective memory and pop-out in extrastriate area V4. *J. Neurosci.* 14, 2190–2199.
- Muller, J.R., Philiastides, M.G., and Newsome, W.T. (2005). Microstimulation of the superior colliculus focuses attention without moving the eyes. *Proc. Natl. Acad. Sci. USA* 102, 524–529.
- Nienborg, H., and 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.
- Nissen, M.J., and Corkin, S. (1985). Effectiveness of attentional cueing in older and younger adults. *J. Gerontol.* 40, 185–191.
- Parker, A.J., and Newsome, W.T. (1998). Sense and the single neuron: probing the physiology of perception. *Annu. Rev. Neurosci.* 21, 227–277.

- Poulet, J.F., and Petersen, C.C. (2008). Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* 454, 881–885.
- Rajan, K., Abbott, L.F., and Sompolinsky, H. (2010). Stimulus-dependent suppression of chaos in recurrent neural networks. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 82, 011903.
- Ray, S., and Maunsell, J.H.R. (2010). Differences in gamma frequencies across visual cortex restrict their possible use in computation. *Neuron* 67, 885–896.
- Reynolds, J.H., and Chelazzi, L. (2004). Attentional modulation of visual processing. *Annu. Rev. Neurosci.* 27, 611–647.
- Reynolds, J.H., and Heeger, D.J. (2009). The normalization model of attention. *Neuron* 61, 168–185.
- Rossi, A.F., and Paradiso, M.A. (1995). Feature-specific effects of selective visual attention. *Vision Res.* 35, 621–634.
- Saenz, M., Buracas, G.T., and Boynton, G.M. (2002). Global effects of feature-based attention in human visual cortex. *Nat. Neurosci.* 5, 631–632.
- Saenz, M., Buracas, G.T., and Boynton, G.M. (2003). Global feature-based attention for motion and color. *Vision Res.* 43, 629–637.
- Sally, S.L., Vidnyansky, Z., and Pappas, T.V. (2009). Feature-based attentional modulation increases with stimulus separation in divided-attention tasks. *Spat. Vis.* 22, 529–553.
- Serences, J.T., Schwarzbach, J., Courtney, S.M., Golay, X., and Yantis, S. (2004). Control of object-based attention in human cortex. *Cereb. Cortex* 14, 1346–1357.
- Shadlen, M.N., Britten, K.H., Newsome, W.T., and Movshon, J.A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* 16, 1486–1510.
- Siegel, M., Donner, T.H., Oostenveld, R., Fries, P., and Engel, A.K. (2008). Neuronal synchronization along the dorsal visual pathway reflects the focus of spatial attention. *Neuron* 60, 709–719.
- Treue, S., and Martinez Trujillo, J.C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399, 575–579.
- Vaadia, E., Haalman, I., Abeles, M., Bergman, H., Prut, Y., Slovin, H., and Aertsen, A. (1995). Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* 373, 515–518.
- van den Heuvel, M.P., and Hulshoff Pol, H.E. (2010). Exploring the brain network: a review on resting-state fMRI functional connectivity. *Eur. Neuropsychopharmacol.* 20, 519–534.
- von Wright, J.M. (1970). On selection in visual immediate memory. *Acta Psychol. (Amst.)* 33, 280–292.
- Wolfe, J.M. (1998). Visual search. In *Attention*, H. Pashler, ed. (London, UK: University College London Press), pp. 13–74.
- Wolfe, J.M., Horowitz, T.S., Kenner, N., Hyle, M., and Vasan, N. (2004). How fast can you change your mind? The speed of top-down guidance in visual search. *Vision Res.* 44, 1411–1426.
- Womelsdorf, T., and Fries, P. (2007). The role of neuronal synchronization in selective attention. *Curr. Opin. Neurobiol.* 17, 154–160.
- Womelsdorf, T., Schoffelen, J.M., Oostenveld, R., Singer, W., Desimone, R., Engel, A.K., and Fries, P. (2007). Modulation of neuronal interactions through neuronal synchronization. *Science* 316, 1609–1612.
- Yantis, S., and Serences, J.T. (2003). Cortical mechanisms of space-based and object-based attentional control. *Curr. Opin. Neurobiol.* 13, 187–193.
- Zohary, E., Shadlen, M.N., and Newsome, W.T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370, 140–143.