nature neuroscience

Attention improves performance primarily by reducing interneuronal correlations

Marlene R Cohen & John H R Maunsell

Visual attention can improve behavioral performance by allowing observers to focus on the important information in a complex scene. Attention also typically increases the firing rates of cortical sensory neurons. Rate increases improve the signal-to-noise ratio of individual neurons, and this improvement has been assumed to underlie attention-related improvements in behavior. We recorded dozens of neurons simultaneously in visual area V4 and found that changes in single neurons accounted for only a small fraction of the improvement in the sensitivity of the population. Instead, over 80% of the attentional improvement in the population signal was caused by decreases in the correlations between the trial-to-trial fluctuations in the responses of pairs of neurons. These results suggest that the representation of sensory information in populations of neurons and the way attention affects the sensitivity of the population may only be understood by considering the interactions between neurons.

The responses of sensory neurons are variable, and laboratory studies typically deal with this variability by averaging responses to many stimulus presentations. In the real world, however, people and animals must respond to individual stimulus events, and the brain is thought to compensate for neuronal variability by encoding sensory information in the responses of large populations of neurons. To understand the way sensory information guides behavior in everyday life, we need to understand the way information is encoded in populations of neurons.

One way to identify the important aspects of a population code is to look at the differences between the neuronal representation of a sensory stimulus when it is used to guide behavior and when it is behaviorally irrelevant. Tasks that control attention provide a powerful way to manipulate behavioral relevance. Attention allows observers to select the most important stimuli and greatly improves perception of the attended location or feature. Attention modulates the firing rates of sensory neurons, typically increasing responses to attended stimuli[1–4]. This increased rate of firing acts to improve the signal-to-noise ratio of individual neurons[5,6] and a recent study found that attention can cause a small additional reduction in the mean-normalized variance (Fano factor) of the responses of some neurons in visual area V4 (ref. 7). However, the net effect of attention on the signal-to-noise ratio of single neurons is modest, suggesting that attention causes large improvements in psychophysical performance by affecting population responses in ways that cannot be measured in single neurons.

Attention could also alter the reliability of neuronal representations by affecting the amount of noise that is shared across a population of neurons. Variability in a population partly depends on the variability of single neurons, but can depend greatly on the extent to which variability is shared across the population. The effect of correlated variability on population sensitivity depends on the way in which the population is read out[8,9], but its effect can be far greater than the effect of independent variability of single neurons. If the noise in individual neurons is independent, averaging the responses of many neurons will lead to a very accurate estimate of the mean, no matter how noisy the individual neurons are. If, however, there are positive correlations in the trial-to-trial fluctuations of the responses of pairs of neurons, then the shared variability can never be averaged out, leading to a more variable (and less accurate) estimate of the mean activity in the population[10–12].

We found that attention adaptively decreased correlated variability in a population of neurons in visual area V4 in a change-detection task. Furthermore, we found that this decrease accounted for the vast majority of the attentional improvement in the amount of sensory information encoded by the population and is probably the major contributor to the improved psychophysical performance. These results indicate that studies that focus on a single neuron, which necessarily ignore interactions between neurons, miss the most critical aspect of the way sensory information is encoded in populations of neurons.

RESULTS

We investigated the effect of attention on both single neuron responses and correlated variability by recording from populations of neurons in V4 using chronically implanted microelectrode arrays in two rhesus monkeys (*Macaca mulatta*). Each monkey had two arrays, allowing us to monitor populations of neurons in both hemispheres simultaneously. The diameter of V4 receptive fields is approximately equal to eccentricity[13,14], so the receptive fields the neurons recorded in a hemisphere typically overlapped at least partially (**Fig. 1a**). We recorded from 376 single units and 2,746 multiunit clusters during 41 d of recording (including 66,578 simultaneously recorded pairs in the same hemisphere and 59,990 pairs in opposite hemispheres).

Howard Hughes Medical Institute and Harvard Medical School Department of Neurobiology, Boston, Massachusetts, USA. Correspondence should be addressed to M.R.C. (marlene_cohen@hms.harvard.edu).

Received 9 July; accepted 30 September; published online 15 November 2009; doi:10.1038/nn.2439

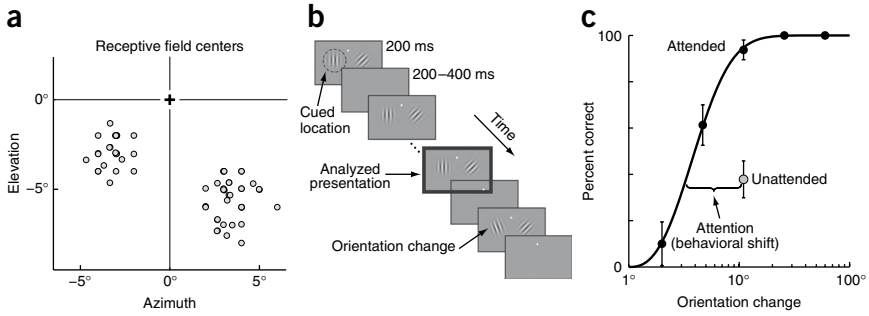

Figure 1 Methods and behavior. (**a**) Center of visual receptive fields for the multiunit signals from one monkey. (**b**) Orientation change–detection task. Two Gabor stimuli synchronously flashed on for 200 ms and off for a randomized 200–400-ms period. At an unsignaled and randomized time, the orientation of one of the stimuli changed, and the monkey was rewarded for making a saccade to the stimulus that changed. Attention was cued in blocks, and the cue was valid on 80% of trials, meaning that on an 'attend-left' block of trials (depicted here), 80% of the orientation changes were to the left stimulus. The monkey was rewarded for correctly detecting any change, even on the unattended side. Unless otherwise stated, all analyses were performed on responses to the stimulus before the orientation change (black outlined panel). (**c**) Psychometric performance from a typical example experiment. The percent correct as a function of orientation change in degrees for trials in which the change occurred at the attended (black points) or unattended (gray point) location is shown. Unattended changes occurred only at the middle difficulty level (11°). Attentional improvement in behavior was quantified as the lateral shift between the percent correct on unattended trials and the fitted psychometric curve for attended trials.

We did not find any important differences between single and multiunits or between the two monkeys, and our population analyses required large neural populations, so we combined single and multiunits (see **Supplementary Results**). However, the statistics that we used for single units are based on a subset of 187 single units that we were confident are unique (if there was a single unit on a given electrode on multiple days, it was only counted once).

The monkeys performed an orientation change–detection task in which spatial attention was manipulated (**Fig. 1b**). Two Gabor stimuli flashed on and off and the monkey's task was to detect a change in the orientation of either stimulus. We manipulated attention in blocks by cueing the monkey as to which stimulus was more likely to change (see Online Methods). Each day, the location, size, orientation and spatial frequency of the Gabors were optimized for a selected single unit in each hemisphere. The two stimuli were therefore different, so directing attention to one of the two stimuli probably modulated both feature-based and spatial attention. Because we recorded from neurons with a wide range of receptive field locations and tuning, most neurons were not well driven by the stimulus (mean driven rate was 8.2 spikes per s for single units and 21.5 spikes per s for multiunits, compared with mean spontaneous rate 5.8 spikes per s for single units and 14.1 spikes per s for multiunits).

Attention greatly improved behavioral performance in this task. To motivate the monkeys to attend to the cued location, we changed the stimulus at the attended location on 80% of trials (trials in which the attentional cue was valid). On the remaining 20% of trials (invalid trials), we tested performance at the unattended location using only a single orientation change (11°), which allowed us obtain reliable estimates of behavior and neural responses despite there being relatively few invalid trials. In an example of a typical recording session, the proportion of trials in which the monkey successfully detected an 11° orientation change was substantially greater on trials when the attended, rather than the unattended, stimulus changed orientation (**Fig. 1c**).

To compute the effects of attention on neural responses during the period in which the monkey's attentional state was most likely to affect its behavioral performance, we focused our analyses on the stimulus presentation directly preceding the orientation change (**Fig. 1b**). On a given day, the stimuli immediately before the orientation change were identical, regardless of the attentional condition, validity of the attentional cue or size of the orientation change. Invalid trials were randomly interleaved with valid trials, so the neuronal effects of attention were indistinguishable on valid and invalid trials. We observed some adaptation of V4 responses between the first and the second stimulus presentation on each trial, but the average responses to the second through tenth stimuli were statistically indistinguishable (t test, $P > 0.5$). Because the orientation change occurred no sooner than the third stimulus presentation, the responses to the stimulus directly before the change were unaffected by the length of the trial.

Consistent with previous studies[3,15–18], we found that attention increased V4 firing rates (**Fig. 2a**). To quantify the increase, we calculated a standard modulation index (MI_{rates}), which was the difference between the average firing rates on trials in which the attended stimulus was inside of or outside of the neuron's receptive field trials divided by the sum (see **Supplementary Results**). The mean MI_{rates} was 0.049 for single units and 0.042 for multiunits, both of which were significantly greater than zero (t test, $P < 10^{-6}$ for single units, $P < 10^{-20}$ for multiunits).

We also found that attention reduces the trial-to-trial variability of individual neurons over a similar time course to its effects on firing rate. As is common to stimulus responses in many cortical areas[19], we observed a drop in the Fano factor (the ratio of the variance of the firing rates to the mean) following stimulus onset (**Fig. 2b**). Following the drop associated with the response transient, the Fano factor remained at a significantly lower level in the attended than in the unattended condition (mean MI_{FF} during the sustained response was −0.011 for single units and −0.017 for multiunits, $P < 0.05$ and $P < 10^{-3}$, respectively). Because the Fano factor (**Fig. 2b**) was calculated using subdistributions of neurons such that the mean firing rates were the same for each time point and attentional condition[19,20], the time course and attentional dependence of the Fano factor were independent of changes in firing rate (see Online Methods).

Attention improved the signal-to-noise ratio of individual V4 neurons (**Fig. 2a**,**b**), but we found that the effect of attention on the correlated variability in pairs of neurons was even more important. For each pair of simultaneously recorded neurons and each attentional condition, we calculated the correlation coefficient between spike count responses to the stimulus preceding the orientation change. This metric, termed noise correlation, measures the correlation in trial-to-trial fluctuations in responses and therefore has a very different timescale than the millisecond timescale synchrony that has been shown to increase with attention[21]. We did not focus on synchrony here because no more pairs exhibited significant synchrony than would be expected by chance (out of 66,578 pairs, 3,609 had significant synchrony in the attended condition and 3,634 had significant synchrony in the unattended condition, 5.4% and 5.5%, respectively;

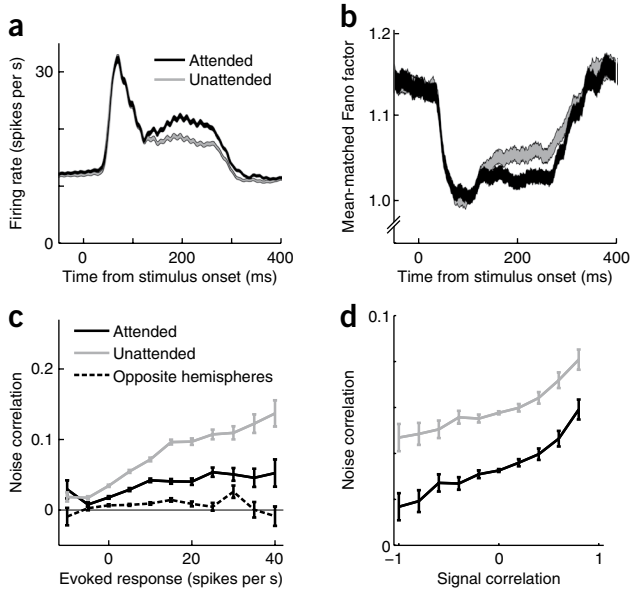

Figure 2 Attentional modulation of firing rate, Fano factor and noise correlation. (**a**) Attention increased firing rates. A peristimulus time histogram of firing rates for all 3,498 single neurons and multiunit clusters on trials in which the stimulus in the same hemifield as the neuron's receptive field was attended (black line) or unattended (gray line) is shown. Line width represents the s.e.m. Ripples reflect the 85-Hz frame rate of the video display. (**b**) Attention decreased mean-matched Fano factor. Plotting conventions are as described in **a**. (**c**) Attention decreased noise correlation. Spike count noise correlation (for responses over the period from 60 to 260 ms following stimulus onset) is plotted as a function of the mean stimulus modulation for the pair of neurons (firing rate during the stimulus–firing rate during the interstimulus blank period). For pairs of neurons in the same hemisphere, correlation was lower when the stimulus in the neurons' receptive field was attended (black line) than when it was unattended (gray line). Pairs of neurons in opposite hemispheres (dashed lines) had correlations that were close to zero. Error bars represent s.e.m. (**d**) Raw noise correlation, but not attentional modulation, depended on signal correlation. Mean noise correlation is plotted as a function of signal correlation, which can be thought of as the similarity in spatial and feature tuning of the two neurons (see Online Methods). As has been previously reported, noise correlation increases with signal correlation. However, the difference in correlation between the attended (black line) and unattended (gray line) conditions did not depend on signal correlation. Error bars represent s.e.m.

$P < 0.05$, bootstrap test described in Online Methods) and synchrony in the attended and unattended conditions was not different (paired t test, $P = 0.46$). Many spikes are needed to detect statistically significant synchrony and even more are needed to detect modulation of synchrony by processes such as attention. Synchrony has therefore been observed in some studies of visual cortex (for example, see refs. 21,22), but not in others[10,23,24] (see ref. 21 for a discussion of the statistical power needed to detect synchrony). The absence of synchrony in our study is probably a result of a combination of the low firing rates of many of our cells caused by stimuli that were suboptimal for most cells, the fact that we calculated synchrony using pairs of spiking neurons rather than correlating spike times with local field potentials[21] and the fact that most neuron pairs were separated by millimeters in the cortex. The correlations that we observed were fluctuations on a longer timescale than millisecond-level synchrony. One possibility is that the same mechanisms that cause low-frequency oscillations in electroencephalography and local field potentials (which have been shown to desynchronize with attention[25–28]) caused the correlations we measured.

To obtain accurate estimates of noise correlation, we did not calculate a time course of correlation as we did for rate and Fano factor because, over short periods, the distributions of spike counts became non-Gaussian (because spike counts can never be negative) and discrete. Skewed, discrete distributions pose a problem for second-order statistics such as correlation, causing noise correlations to approach zero as the mean number of spikes decreases[22,29,30]. We therefore calculated noise correlation over the entire 200-ms interval (**Fig. 2c**).

Because the stimuli produced a wide range of responses across the population of neurons, we binned the neuron pairs by their mean evoked response across both attentional conditions (driven rate–baseline; **Fig. 2c**). Noise correlations were highest for pairs of neurons in the same hemisphere that both responded strongly to the stimulus. This result can be explained by the fact that correlations tend to increase with firing rate[30] and the observation that noise correlations are highest for neurons with similar tuning[10,22,24,29]. Our dataset included neurons with a broad range of preferences for orientation and other stimulus properties and different receptive field locations, so two neurons that were both strongly modulated by the stimulus likely had similar tuning.

To test the effect of tuning similarity on noise correlation more directly, we calculated noise correlation as a function of signal correlation (**Fig. 2d**). In a separate set of trials, we presented Gabor stimuli at a variety of locations and orientations while the monkey performed a change-detection task far outside the neurons' receptive fields (see Online Methods). We calculated signal correlation by computing a correlation between the mean responses of each neuron to each stimulus. Consistent with previous results[10,22–24], we found that noise correlation is highest for neurons with similar tuning (large, positive signal correlation) and lowest for neurons with opposite tuning (negative signal correlation). Unlike a recent study of noise correlations in V1 using the same electrode arrays that we used here[29], we found that noise correlation did not depend on cortical distance. We suspect that the greater retinotopic and tuning organization of V1 compared with V4 accounts for the differences in our results.

We found that even for the least responsive neurons (**Fig. 2c**) and pairs of neurons with dissimilar stimulus preferences (**Fig. 2d**), correlations within a hemisphere were on average positive, indicating that there is shared variability throughout the population. In contrast, we found that noise correlations for pairs of neurons in opposite hemispheres were close to zero (**Fig. 2c**), meaning that trial-to-trial fluctuations in the two hemispheres are independent within an attentional condition.

The biggest physiological effect of attention in our dataset was a large decrease in the correlations between pairs of neurons in the same hemisphere (**Fig. 2c,d**). On average, attention reduced noise correlations by about half (mean $MI_{cor} = -0.35$ for single units and -0.29 for multiunits, $P < 10^{-5}$ for single units and $P < 10^{-9}$ for multiunits). Attentional modulation of correlation depended strongly on how much the neurons were driven by the stimulus; for the most responsive pairs of neurons, noise correlation in the attended condition was roughly one-third the correlation in the unattended condition (**Fig. 2c**). In contrast, the effect of attention on correlations did not depend on the degree of tuning similarity between the two cells (**Fig. 2d**). This observed decrease in correlation as a result of attention is the opposite result predicted by the mathematical relationship between firing rate and correlation[30]. Attention tends to increase firing rates (**Fig. 2a**), which makes the distributions of spike counts more Gaussian and less discretized, leading to a predicted increase

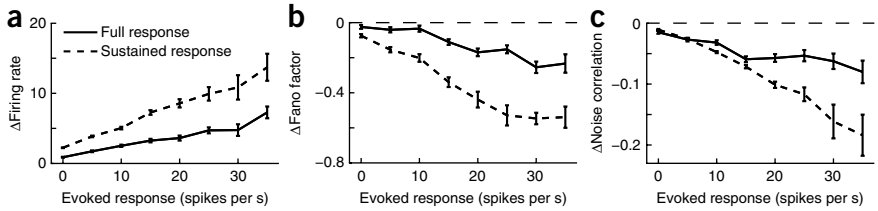

Figure 3 Attention has the biggest effects on the most responsive neurons. (**a**) Difference in mean firing rate between trials when the stimulus in the neuron's receptive field was attended and unattended as a function of stimulus modulation (rate during stimulus period–interstimulus period) during the full stimulus period (solid line) or the sustained response following the onset transient (dashed line). Error bars represent s.e.m. (**b**) Data are presented as in **a** for Fano factor. (**c**) Data are presented as in **a** for noise correlation for pairs of neurons in the same hemisphere.

in correlation. Therefore, decreases in correlation cannot be a simple mathematical consequence of increases in firing rate.

Previous studies have shown that attention modulates firing rate more for neurons with the biggest response to the stimulus (whose responses may be more informative for the task)[15–18]. Furthermore, a recent study found that fast-spiking neurons with high firing rates (putative interneurons, separated from putative excitatory neurons on the basis of waveform width) had larger differences in Fano factor than regular-spiking neurons[7]. Our hardware filters prevented us from distinguishing these neuron types on the basis of waveform, but consistent with this study and studies of attentional modulation of firing rate, we found that attention had a bigger effect on the rates, Fano factors and noise correlations of neurons that responded strongly to the stimulus (**Fig. 3**). Neurons (or pairs of neurons) that were most strongly driven by the stimulus (biggest difference between evoked and baseline firing rate) probably have receptive field locations and tuning properties that make them well suited for this task, and these neurons showed the largest effects of attention by all three measures.

Recording from both hemispheres simultaneously allows us to be sure that the correlation changes that we observed were spatially specific effects of attention. The same block of trials that yielded low correlations in one hemisphere gave high correlations in the other, so nonspecific factors such as arousal or motivation cannot account for the changes in correlation that we observed. The fact that trial-to-trial variability in the two hemispheres was virtually independent is further evidence that the correlation changes we observed in a hemisphere are spatially specific.

Consistent with many previous studies (for examples, see refs. 21,31–34), we found that attention primarily affected the sustained part of the response rather than the onset transient (**Fig. 2a**,**b**). In our data, attentional modulation of firing rate became statistically significant 122 ms after stimulus onset (the first time point at which the 95% confidence intervals for the means of the two attentional conditions did not overlap). In addition to examining the effect of attention on rates, Fano factor and correlations during the entire stimulus period, we calculated the attentional effects for all three measures during the sustained response of the response (122–260 ms after stimulus onset). As expected, attentional effects were larger during the sustained period by what appeared to be a fairly constant factor (**Fig. 3**).

Attention changed the responses of both single neurons and correlated variability in ways that could allow each to contribute to improvements in population sensitivity (**Figs. 2** and **3**). A primary goal of our study was to determine the relative importance of changes in firing rates, Fano factor and noise correlations. Because Fano factor measures the variability of single neurons without regard to the source of that variability, the decrease in Fano factor that we observed (**Figs. 2b** and **3b**) could arise from a decrease in the independent variability of individual neurons, a decrease in shared variability across the population or a combination of both. Noise correlation measures the degree of shared variability. We therefore focused on the other aspect of variability captured by the Fano factor and asked how much of a decrease in independent variability that was large enough to account for the full decrease in Fano factor would improve population sensitivity.

We compared the effects of attentional modulation of firing rates, independent variability and noise correlation and found that the modulation of correlation had by far the greatest influence on the attentional improvement in population sensitivity (**Fig. 4**). We first quantified the degree to which attention improved the sensitivity of the groups of neurons that we recorded (**Fig. 4a**,**c**) and then determined the amount of that improvement that was caused by attentional modulation that affected only rate, only independent variability or only correlation (**Fig. 4d**,**e**).

We quantified how much attention improved neuronal signals in our recorded populations (schematized for a hypothetical two-neuron dataset in **Fig. 4a** and shown for a real 38-neuron dataset in **Fig. 4c**). The monkey's task was to detect a change in the orientation of the stimulus, so we defined population sensitivity as the discriminability between the distributions of responses to the original orientation and the changed orientation. For each of the single and multi-units that we recorded from a given hemisphere on a given day (mean of 39.5 neurons, range of 14–74), we calculated responses to the stimulus preceding the change from 60 ms to 260 ms following stimulus onset and the changed stimulus starting 60 ms after onset and continuing for either 200 ms or until 60 ms before the monkey's response, whichever came first. The mean time from the onset of the changed stimulus to the onset of the monkey's response was 251 ms, and 260 ms fell at least 60 ms before the saccade on 39% of trials. We experimented with other intervals for computing spike counts (including identical periods for the original and changed stimuli (200 ms each) and also cutting off the response to the changed stimulus 100 ms or 0 ms before the saccade) and these did not qualitatively affect the proportion of the improvement in population sensitivity accounted for by each of the three factors that we considered. Using this time period, attentional modulation during the changed stimulus was indistinguishable from modulation during the previous stimulus (**Supplementary Fig. 1**).

We plotted one point for each stimulus in each trial in an n-dimensional space in which each dimension corresponds to the response of one of the n neurons that we recorded in a given hemisphere (**Fig. 4a**). We then calculated the mean response for each stimulus and projected all responses onto an 'axis of discrimination' drawn through the two means. This was done separately for the two attention conditions, producing pairs of one-dimensional distributions of projections for each attention condition (**Fig. 4c**).

We measured population sensitivity by calculating d' (the difference in the means divided by their root mean square s.d.), which is monotonically related to theoretical performance on classifying stimuli, so attention should increase d' to improve behavioral performance. We quantified the attentional improvement in population sensitivity as the difference in d' between the attended and unattended conditions. We then normalized the d' values to reflect the measured improvement. The amount of attentional improvement in our d' measure correlated

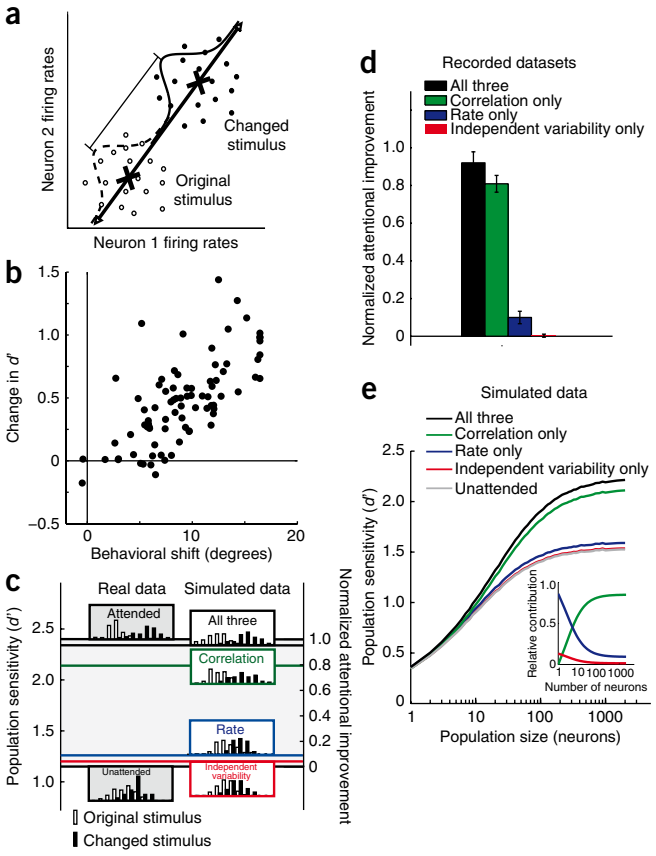

Figure 4 Modulation of noise correlation accounts for the majority of the attentional improvement in population sensitivity. (**a**). Procedure for calculating the sensitivity of the population. For each trial and attentional condition, the firing rate response of the *n* neurons recorded simultaneously in a given hemisphere to the stimulus immediately before the orientation change (open circles) and the changed stimulus (filled circles) is plotted as a point in an *n*-dimensional space (a fictional two-neuron example is plotted here). Each point is projected onto the axis connecting the center of mass of the cloud of points for each time period (X's), leaving a one-dimensional distribution of projected values for each time period (dashed and solid curves). The sensitivity of the population to the change in the stimulus is quantified as the discriminability of the two distributions in units of *d′*. (**b**). Population *d′* and behavioral improvement were highly correlated. For each hemisphere day, population *d′* is plotted as a function of the behavioral improvement (quantified as the lateral shift between performance at the unattended location and the fitted psychometric curve for the attended condition). (**c**). Procedure for calculating the amount of the observed attentional improvement explained by each factor for a representative example dataset. Histograms of projections onto the axis defined in **a** are plotted for the real data (left column, for attended and unattended trials) and for simulations (right column). We defined the observed attentional improvement as the difference between *d′* for the two attentional conditions (*d′* = 2.40 for the attended condition and 1.15 for the unattended condition, giving an improvement of 1.25 in this example). The left axis represents *d′* and the right axis represents normalized proportion of attentional improvement (by definition 1.0 for the attended condition and 0.0 for the unattended condition). We calculated the fraction of the observed attentional improvement explained by each factor(s) by comparing the simulated *d′* (right column of distributions, see Online Methods) to the *d′* for the real unattended data. (**d**). Average proportion of actual attentional improvement for all 82 datasets (one dataset for each hemisphere day). Error bars represent s.e.m. (**e**). Population sensitivity as a function of the number of neurons involved in the task. Population *d′* was calculated using the method described in **a** and **b**, except that data in both the attended and the unattended conditions were simulated. For each population size, we sampled, with replacement, from the entire population of neurons from all datasets combined. Each simulation was run 100 times for 10,000 trials on each run. The inset plots the relative contribution of each factor (which is the ratio of the improvement in *d′* for that factor alone to the improvement in *d′* for all three factors) as a function of population size.

strongly with the monkey's behavioral improvement resulting from attention. For each hemisphere day, we quantified behavioral improvement as the lateral shift between measured performance in the unattended condition and the fitted psychometric curve in the attended condition (**Fig. 1c** and Online Methods). Attention shifted the psychometric curve by 7.7°, which was typical for our datasets (7.6° ± 0.5°, mean ± s.e.m.; **Fig. 1c**). Attentional improvement in neuronal *d′* (attended-unattended) for each hemisphere day was highly correlated with behavioral improvement ($R = 0.69$, $P < 10^{-12}$; **Fig. 4c**). This strong correlation suggests that our *d′* metric captures the important aspects of the improvements in population sensitivity that lead to improvements in behavior.

Each of the physiological changes that we observed in rate, Fano factor and correlation could have contributed to the improvement in population sensitivity. We next compared how much attentional modulation of each factor alone and the three factors together contributed to the actual improvements in *d′* that we calculated. To isolate the contribution of each factor, we simulated the responses of populations of neurons using the same mean rates, noise correlations, Fano factors and number of neurons as the groups of neurons that we recorded (see Online Methods) and compared the calculated *d′* for each simulation to the real data in the unattended condition. In this example dataset, attentional modulation of all three factors together (**Fig. 4c**) accounted for 95% of the attentional improvement that we observed in the real data. We then calculated the contribution of each factor separately by simulating attentional modulation of the factor of interest and using the values observed in the unattended condition for the other two factors (**Fig. 4c**). Correlation alone accounted for 79% of the attentional improvement, rate accounted for 9% and modulation of independent variability accounted for 4%.

This example is typical of the 82 datasets (**Fig. 4c**). On average, attentional modulation of the three factors together accounted for 92% of the attentional improvement that we observed in the actual populations (**Fig. 4d**). Notably, this result means that population sensitivity is well modeled by accounting only for rate, independent variability and pair-wise noise correlation, and that any other factors (including any higher-order correlations) account for no more than 8% of the observed improvement in population sensitivity. Consistent with this, population responses in the retina are well described by the responses of individual neurons and pair-wise correlations[35,36]. Overall, modulation of noise correlation was by far the most important factor in explaining the improvement in population sensitivity. Attentional modulation of noise correlation accounted for 81% of the observed improvement, rate accounted for 10% and independent variability accounted for only 0.3% (which was not significantly different than 0.0, *t* test, $P = 0.82$).

Unsurprisingly, we found that both the observed raw population *d′* and the improvement in *d′* resulting from attention depended on the number of neurons that we recorded. Because there is no a priori way of knowing how many neurons are involved in the task, we examined the dependence of these measures on population size by sampling, with replacement, the firing rates, Fano factors and correlations of all

of the neurons that we recorded over all recording sessions (see Online Methods and refs. 11,24). We calculated population d' for simulations in which attention modulated either all three factors, one factor individually or none of the factors (**Fig. 4e**). In all cases, d' increased with population size. Because noise correlations are on average positive for both attentional conditions, d' asymptotes for large populations[9–11]. Modulation of noise correlation accounted for most of the attentional improvement in sensitivity for population sizes greater than five neurons (**Fig. 4e**). For very small populations, however, this was necessarily not true (**Fig. 4e**). If performance depends on a single neuron, there can be no correlation, and the small attentional improvement depends almost entirely on modulation of firing rate. In our simulations, noise correlations became dominant for populations of more than five neurons. If anything, this estimate may be high because we recorded from many neurons with stimulus preferences that were not well matched to the stimuli that we presented, resulting in low firing-rate responses and low noise correlations (**Fig. 2c**). If we had recorded from neurons that were better matched to the stimuli, correlations would probably have been higher (**Fig. 2c**), shifting the point at which correlation becomes most important to population sizes even lower than five neurons. Many more than five neurons are thought to be involved in virtually every task, so changes in correlation probably dominate attentional improvement in nearly all situations.

DISCUSSION

Why do changes in shared variability have a bigger effect on population sensitivity than changes in the signal-to-noise ratio of single neurons? One answer is that the changes in correlated variability that we observed were larger than the changes in firing rate or Fano factor. However, we re-ran the simulations (**Fig. 4e**), assuming that the three factors all had the same modulation index as the changes in rate (see **Supplementary Results**), and correlation still dominated for population sizes greater than 30 neurons. Instead, the explanation lies in the fact that no matter how noisy individual neurons are, independent variability can be averaged out if the population of neurons is large enough. Correlated variability, however, can never be averaged out by simply adding neurons to the population.

Noise correlation can either improve or reduce population sensitivity, depending on the algorithm by which neural responses are read out[8,9], and our simulations could, in principle, have revealed that the observed correlation decreases acted to reduce population sensitivity. However, theoretical studies have shown that decreased correlation improves discrimination if the difference between the responses to the stimuli to be discriminated (the original and changed stimuli in our task) are of the same sign for most neurons[8,9], which turned out to be the case in our dataset.

Most of the neurons that we recorded (92%) responded more strongly to the changed than the previous stimulus, presumably reflecting adaptation to the series of identical stimuli preceding the change. Therefore, the optimal quantity to be read out is similar to a (positively) weighted mean of the responses of the population and the axis of discrimination that we determined (**Fig. 4a**) was close to the weighted population mean. The attention-related decrease in correlation therefore improved the sensitivity of the population by reducing the amount of shared variability that could not be removed by averaging. In contrast, a recent study found no effect of attention on noise correlations in a situation in which correlations were shown to have no effect on the sensitivity of the population[37]. There are further situations (such as those in which the optimal readout algorithm is more similar to a subtraction of two populations of cells) in which an increase in correlation would improve the sensitivity of the population[8,9]. Whether attention would increase correlations in such tasks remains to be determined.

It is likely, of course, that the brain uses a different algorithm for extracting stimulus information from the responses of many neurons than the very simple decoding scheme that we used (**Fig. 4**). However, the observation that attentional modulation of noise correlation explains most of the attentional improvement in population sensitivity is probably true for any sensible decoding algorithm. First, the difference in the amount of attentional improvement explained by pair-wise correlations was very large compared with the amount explained by the changes in the responses of single neurons, suggesting that noise correlations will dominate using any decoding algorithm. Furthermore, correlation was by far the most important factor using any of several linear discriminators that we tried, including the single axis projection described here, Fisher discriminants and support vector machines (data not shown). Higher-order decoders that explicitly read out interneuronal correlations[38–41] will be even more affected by attentional modulation of correlation than linear discriminators. Finally, any sort of decoding algorithm that incorporates a mean (or weighted mean) of the responses of many neurons will be greatly affected by noise correlations[10–12].

Mathematically, correlation is invariant to the mean response (the correlation coefficient is the ratio of the covariance to the square root of the product of the individual variances, so both the numerator and denominator are proportional to the product of the means), so underlying noise correlations cannot be changed by a simple scaling of neural responses (that is, a gain change). Instead, noise correlations in cortex are thought to arise primarily from common, noisy inputs[10,22,23,29]. The fact that attention primarily decreased correlations provides clues about the mechanisms by which attention affects populations of sensory neurons. A decrease in correlation combined with an increase in firing rates is consistent with a decrease in the strength of an effectively inhibitory input that is common across the population. One possibility is that attention results in a decrease in the weights or activity of inputs that cause divisive normalization, a mechanism that normalizes responses to many stimuli in a receptive field and has recently been proposed to underlie attention[34,42,43]. In fact, we found a correlation between the mean attentional modulation of the firing rates of a pair of neurons and modulation of their noise correlation ($R = -0.32$, $P < 10^{-4}$) and also between the average rate and correlation changes in a hemisphere day ($R = -0.61$, $P < 10^{-9}$; see **Supplementary Fig. 2**), which is consistent with the idea that the two attentional changes may be mediated by the same mechanism.

Attention improves perception of the attended location or feature, so studying the effects of attention on populations of sensory neurons reveals the aspects of the population code that are most important for accurately encoding information about a behaviorally relevant stimulus. Here we found that attention improved population sensitivity primarily by changing noise correlations and even the small pair-wise correlations that we observed had a marked effect on the sensitivity of the population. Therefore, understanding the interactions between pairs of neurons is critical for understanding population coding (see also refs. 8–11,38–41,44).

Rather than examining mean responses over many trials, the brain makes decisions on the basis of the responses of many neurons over a short period. Our results indicate that studies of average responses of single neurons miss interactions between neurons that have critical effects on behavior. Together, these results suggest that the future of studying population coding will rely on multi-electrode or

imaging technologies that allow glimpses of population coding on the timescale of a single behavioral decision.

Note added in proof: A recent study in area V4 confirmed that attention reduces noise correlations[45].

METHODS

Methods and any associated references are available in the online version of the paper at http://www.nature.com/natureneuroscience/.

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

ACKNOWLEDGMENTS

We thank D. Ruff for spike sorting and monkey training assistance, M. Churchland for code and advice regarding mean-matching and Fano factor analysis, and M. Histed for many helpful discussions. They, A. Ni and A. Smolyanskaya provided comments on an earlier version of the manuscript. This work was supported by US National Institutes of Health grant R01EY005911 and the Howard Hughes Medical Institute.

AUTHOR CONTRIBUTIONS

M.R.C. collected the data and performed the analyses. M.R.C. and J.H.R.M. designed the study and wrote the paper.

Published online at http://www.nature.com/natureneuroscience/.
Reprints and permissions information is available online at http://www.nature.com/reprintsandpermissions/.

1. Yantis, S. & Serences, J.T. Cortical mechanisms of space-based and object-based attentional control. *Curr. Opin. Neurobiol.* **13**, 187–193 (2003).
2. Assad, J.A. Neural coding of behavioral relevance in parietal cortex. *Curr. Opin. Neurobiol.* **13**, 194–197 (2003).
3. Reynolds, J.H. & Chelazzi, L. Attentional modulation of visual processing. *Annu. Rev. Neurosci.* **27**, 611–647 (2004).
4. Maunsell, J.H. & Treue, S. Feature-based attention in visual cortex. *Trends Neurosci.* **29**, 317–322 (2006).
5. Tolhurst, D.J., Movshon, J.A. & Dean, A.F. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* **23**, 775–785 (1983).
6. McAdams, C.J. & Maunsell, J.H. Effects of attention on the reliability of individual neurons in monkey visual cortex. *Neuron* **23**, 765–773 (1999).
7. Mitchell, J.F., Sundberg, K.A. & Reynolds, J.H. Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* **55**, 131–141 (2007).
8. Abbott, L.F. & Dayan, P. The effect of correlated variability on the accuracy of a population code. *Neural Comput.* **11**, 91–101 (1999).
9. Averbeck, B.B., Latham, P.E. & Pouget, A. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, 358–366 (2006).
10. Zohary, E., Shadlen, M.N. & Newsome, W.T. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* **370**, 140–143 (1994).
11. Shadlen, M.N., Britten, K.H., Newsome, W.T. & Movshon, J.A. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* **16**, 1486–1510 (1996).
12. Shadlen, M.N. & Newsome, W.T. The variable discharge of cortical neurons: implications for connectivity, computation and information coding. *J. Neurosci.* **18**, 3870–3896 (1998).
13. Desimone, R. & Schein, S.J. Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form. *J. Neurophysiol.* **57**, 835–868 (1987).
14. Gattass, R., Sousa, A.P. & Gross, C.G. Visuotopic organization and extent of V3 and V4 of the macaque. *J. Neurosci.* **8**, 1831–1845 (1988).
15. Motter, B.C. Focal attention produces spatially selective processing in visual cortical areas V1, V2 and V4 in the presence of competing stimuli. *J. Neurophysiol.* **70**, 909–919 (1993).
16. McAdams, C.J. & Maunsell, J.H. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.* **19**, 431–441 (1999).
17. Williford, T. & Maunsell, J.H. Effects of spatial attention on contrast response functions in macaque area V4. *J. Neurophysiol.* **96**, 40–54 (2006).
18. Treue, S. & Martinez Trujillo, J.C. Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* **399**, 575–579 (1999).
19. Churchland, M.M. *et al.* Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Front. Syst. Neurosci.: Comput. Syst. Neurosci. Abstr.* doi:10.3389/conf.neuro.06.2009.03.295 (2009).
20. Churchland, M.M., Yu, B.M., Sahani, M. & Shenoy, K.V. Techniques for extracting single-trial activity patterns from large-scale neural recordings. *Curr. Opin. Neurobiol.* **17**, 609–618 (2007).
21. Fries, P., Reynolds, J.H., Rorie, A.E. & Desimone, R. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* **291**, 1560–1563 (2001).
22. Kohn, A. & Smith, M.A. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* **25**, 3661–3673 (2005).
23. Bair, W., Zohary, E. & Newsome, W.T. Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J. Neurosci.* **21**, 1676–1697 (2001).
24. Cohen, M.R. & Newsome, W.T. Context-dependent changes in functional circuitry in visual area MT. *Neuron* **60**, 162–173 (2008).
25. Thut, G., Nietzel, A., Brandt, S.A. & Pascual-Leone, A. Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *J. Neurosci.* **26**, 9494–9502 (2006).
26. Foxe, J.J., Simpson, G.V. & Ahlfors, S.P. Parieto-occipital approximately 10 Hz activity reflects anticipatory state of visual attention mechanisms. *Neuroreport* **9**, 3929–3933 (1998).
27. Babiloni, C. *et al.* Subsecond 'temporal attention' modulates alpha rhythms. A high-resolution EEG study. *Brain Res. Cogn. Brain Res.* **19**, 259–268 (2004).
28. Bastiaansen, M.C., Bocker, K.B., Brunia, C.H., de Munck, J.C. & Spekreijse, H. Event-related desynchronization during anticipatory attention for an upcoming stimulus: a comparative EEG/MEG study. *Clin. Neurophysiol.* **112**, 393–403 (2001).
29. Smith, M.A. & Kohn, A. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* **28**, 12591–12603 (2008).
30. de la Rocha, J., Doiron, B., Shea-Brown, E., Josic, K. & Reyes, A. Correlation between neural spike trains increases with firing rate. *Nature* **448**, 802–806 (2007).
31. Reynolds, J.H., Pasternak, T. & Desimone, R. Attention increases sensitivity of V4 neurons. *Neuron* **26**, 703–714 (2000).
32. Roelfsema, P.R. & Spekreijse, H. The representation of erroneously perceived stimuli in the primary visual cortex. *Neuron* **31**, 853–863 (2001).
33. Sundberg, K.A., Mitchell, J.F. & Reynolds, J.H. Spatial attention modulates center-surround interactions in macaque visual area v4. *Neuron* **61**, 952–963 (2009).
34. Lee, J. & Maunsell, J.H. A normalization model of attentional modulation of single unit responses. *PLoS One* **4**, e4651 (2009).
35. Shlens, J. *et al.* The structure of multi-neuron firing patterns in primate retina. *J. Neurosci.* **26**, 8254–8266 (2006).
36. Schneidman, E., Berry, M.J., II, Segev, R. & Bialek, W. Weak pair-wise correlations imply strongly correlated network states in a neural population. *Nature* **440**, 1007–1012 (2006).
37. Poort, J. & Roelfsema, P.R. Noise correlations have little influence on the coding of selective attention in area V1. *Cereb. Cortex* **19**, 543–553 (2009).
38. Seriès, P., Latham, P.E. & Pouget, A. Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. *Nat. Neurosci.* **7**, 1129–1135 (2004).
39. Beck, J.M. *et al.* Probabilistic population codes for Bayesian decision making. *Neuron* **60**, 1142–1152 (2008).
40. Pouget, A. & DeAngelis, G.C. Paying attention to correlated neural activity. *Nat. Neurosci.* **11**, 1371–1372 (2008).
41. Pillow, J.W. *et al.* Spatio-temporal correlations and visual signaling in a complete neuronal population. *Nature* **454**, 995–999 (2008).
42. Reynolds, J.H. & Heeger, D.J. The normalization model of attention. *Neuron* **61**, 168–185 (2009).
43. Boynton, G.M. A framework for describing the effects of attention on visual responses. *Vision Res.* **49**, 1129–1143 (2009).
44. Kohn, A., Zandvakili, A. & Smith, M.A. Correlations and brain states: from electrophysiology to functional imaging. *Curr. Opin. Neurobiol.* **19**, 434–438 (2009).
45. Mitchell, J.F., Sundberg, K.A. & Reynolds, J.H. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* **63**, 879–438 (2009).

ONLINE METHODS

The subjects in this experiment were two adult male rhesus monkeys (*Macaca mulatta*, 9 and 12 kg). All animal procedures were in accordance with the Institutional Animal Care and Use Committee of Harvard Medical School. Before training, each monkey was implanted with a head post and a scleral search coil for monitoring eye movements. After the monkey 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 adjacent electrodes was 400 μm. The two arrays were connected to a percutaneous connector that allowed electrophysiological recordings.

We placed the arrays between the lunate and superior temporal sulci, which were visible during surgery. The center of the receptive field for the multiunit signal from each functional electrode from Monkey 1 is shown in **Figure 1a**. Monkey 2 underwent an unscheduled 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 receptive fields than Monkey 1. Other than the receptive field distributions, the main physiological results were indistinguishable for the two monkeys.

The data presented here are from 41 d of recording (20 from Monkey 1 and 21 from Monkey 2), each comprising at least four blocks of each attentional condition (125 successfully completed trials per block). We recorded useful data from 376 unique single neurons (192 from Monkey 1 and 184 from Monkey 2) and 2,746 multiunit clusters (1,070 from Monkey 1 and 1,676 from Monkey 2). All spike sorting was done manually offline using commercial spike-sorting software (Plexon). The dataset consisted of 66,578 simultaneously recorded pairs of neurons (single and multiunits combined) in the same hemisphere and 59,990 pairs in opposite hemispheres. **Figure 4b**,**c** is based on populations of neurons recorded simultaneously in a single hemisphere.

The monkeys performed an orientation change–detection task. The trial began when the monkey fixed a small spot in a 1.5° square fixation window in the center of a video display (85-Hz refresh rate, $1{,}024 \times 768$ pixels, gamma corrected). Two achromatic Gabor stimuli whose size, location, spatial frequency and orientation were each optimized for one neuron in each hemisphere (new optimized neurons and stimuli each day) flashed on for 200 ms and off for a randomized period (200–400 ms picked from a uniform distribution between each stimulus presentation). At an unsignaled time picked from an exponential distribution (minimum of 1,000 ms, mean of 3,000 ms and maximum of 5,000 ms), the orientation of one of the stimuli changed. The monkey was given a liquid reward for making a saccade to the stimulus that changed within 500 ms of its appearance. To account for saccadic latency and to avoid rewarding the monkey for guessing, we rewarded the monkey only for saccades beginning at least 100 ms after the change. If no change occurred within the maximum 5,000 ms, the monkey was rewarded simply for maintaining fixation. Attention was cued to one stimulus location or the other in blocks of 125 trials. Before the start of each block, the monkey performed ten instruction trials (which were not included in any of the analyses presented here) in which there was only a single stimulus. In the upcoming block of trials, 80% of the orientation changes occurred on the same side as the stimulus in the instruction trials, meaning that on an 'attend-left' block of trials, 80% of the orientation changes were to the left stimulus. Only one change occurred on each trial and the monkey was rewarded for correctly detecting any change, even on the unattended side.

Unless otherwise stated, all analyses were performed on responses to the stimulus presentation immediately before the orientation change. The mean-matching procedure for Fano factor (**Fig. 2b**) is described in detail elsewhere[19,20]. Briefly, the mean spike count and variance (and thereby the Fano factor) were calculated for each neuron, attentional condition and 20-ms time interval. The goal was to have the same distribution of mean firing rates (but not variances) at each time point and attentional condition, so we used a different subdistribution of neurons at each time point and condition. We compared distributions of means at each time point and condition and selected the greatest common distribution. We then subsampled our neurons at each time point and condition to match that distribution and then plotted average Fano factor (ratio of the variance to the mean) for those subdistributions.

The analyses in **Figures 2c**, **3** and **4** are based on spike count stimulus responses calculated from the period between 60 and 260 ms after stimulus onset. For the analyses in **Figure 4**, firing rates to the changed stimulus were obtained for spikes from 60 ms after stimulus onset to 260 ms after onset or 60 ms before the saccade, whichever came first. All analyses used only correctly completed trials. The analyses in **Figure 4** are based only on trials that had an 11° orientation change because there were responses to the changed stimulus in both the attended and unattended condition for this difficulty level (see **Fig. 1c**).

Tuning and signal correlation. To assess the tuning of the V4 neurons that we recorded, we measured responses to a variety of Gabor stimuli either before or after the primary experiments on a given day. As the monkey 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 being studied), we synchronously flashed two additional Gabor stimuli in the lower visual field (one in the left and one in the right hemifield) for 100 ms each. The test Gabor stimuli were the same size and spatial frequency as the Gabor stimuli in the main attention task that day. We varied the azimuth (five locations per hemifield), elevation (eight locations) and orientation (six orientations) of the test Gabor stimuli (at least ten repetitions per unique stimulus). We obtained spike count responses during the period from 60–160 ms following stimulus onset. To minimize any effects of adaptation, we only analyzed responses to the stimuli that occurred after the first stimulus and before the changed stimulus in the orientation change task. To compute signal correlation (**Fig. 2d**), we computed the average response of each neuron to each of the 240 unique test Gabor stimuli and computed a correlation coefficient between the average responses.

Simulation of population responses. The analyses in **Figure 4** required us to simulate the responses of populations of V4 neurons whose properties were identical to the ones that we recorded except that attention modulated only one physiological factor (independent variability, firing rate or pair-wise noise correlation) at a time (**Fig. 4b**–**d**) and the number of neurons in the population was varied (**Fig. 4d**). We used methods similar to those described previously[24,11] to impose correlations on simulated populations of neurons. We measured the mean firing rate, Fano factor (as an upper bound for independent variability) and noise correlation for each neuron or pair of neurons and attentional condition. We then simulated responses of neurons with Gaussian distributions of firing rates and the same mean rates and Fano factors as the neurons that we recorded. We imposed correlations in the trial-to-trial fluctuations in responses to match our recorded correlation structure.

To isolate the contribution of attentional modulation of the three physiological factors, we used the measured values of the isolated factor in the attended condition and the other factors in the unattended condition. For example, to isolate the contribution of attentional modulation of firing rate, we used the measured firing rates from the attended condition and the measured Fano factors and noise correlations from the unattended condition. For the simulations in **Figure 4b**,**c**, we simulated responses to the average number of trials the monkey performed in the attended and unattended conditions in that dataset.

To vary population size in **Figure 4d**, we sampled, with replacement, from the entire population of neurons that we recorded in all datasets. For pairs of neurons that were not recorded simultaneously or when we resampled the same neuron more than once in a simulated population, we simulated noise correlation as the average noise correlation for pairs of neurons with a given mean firing rate (**Fig. 2c**). For each population size, we resampled the full set of neurons 1,000 times and simulated responses on 10,000 trials per population.

Synchrony. In addition to noise correlation, which primarily measures correlations on the timescale of tens of milliseconds, we tested for attentional modulation of millisecond-timescale synchrony. We first determined whether each of our 66,578 simultaneously recorded pairs of neurons in a hemisphere exhibited significant synchrony by comparing the measured cross-correlogram (CCG) to a shuffled CCG in each attentional condition. We computed CCGs from the responses during a 60–260-ms time period following the onset of the stimulus preceding the orientation change. To compute the shuffled CCGs, we randomized the trial order for each neuron and then calculated the mean and 95% confidence interval for these shuffled CCGs (1,000 reshuffles). We found that only 3,609 pairs (5.4%) showed significant synchrony from −3 to +3 ms in the attended condition (the integral from −3 to +3 ms of the measured CCG fell outside the 95% confidence interval for the shuffled CCGs) and only 3,634 pairs (5.5%) showed significant synchrony in the unattended condition, which is close to the number expected by chance. We performed a paired *t* test on the differences between the measured integrals in the two conditions for each pair and found that the distributions for the two attention conditions were statistically indistinguishable.