

Influence of On-Off dynamics and selective attention on the spatial pattern of correlated variability in neocortex

Yan-Liang Shi¹, Nicholas A. Steinmetz², Tirin Moore³, Kwabena Boahen⁴, and Tatiana A. Engel^{1,†}

¹*Cold Spring Harbor Laboratory, Cold Spring Harbor, NY*

²*Department of Biological Structure, University of Washington, Seattle, WA*

³*Department of Neurobiology and Howard Hughes Medical Institute, Stanford University, Stanford, CA*

⁴*Departments of Bioengineering and Electrical Engineering, Stanford University, Stanford, CA*

†*Corresponding author e-mail: engel@cshl.edu*

ABSTRACT

Correlated activity fluctuations in neocortex influence sensory responses and behavior. Neural correlations reflect anatomical connectivity and change dynamically with cognitive states, such as attention. Yet, how anatomical connectivity and cognitive states define the population structure of correlations is not known. We measured correlations in single cortical columns and found that the magnitude of correlations, their attentional modulation and dependence on lateral distance are predicted by On-Off dynamics, synchronous fluctuations in population activity across cortical layers. We developed a network model, in which spatial connectivity correlates the On-Off dynamics across nearby columns. We show that attentional inputs modulate the spatial extent of On-Off dynamics, resulting in spatially non-uniform changes in correlations. We confirm this prediction in our columnar recordings by showing that attentional modulation of correlations depends on lateral distance. Our results reveal how heterogeneous spatial patterns of correlations arise from the connectivity and network dynamics during attention.

Neocortical circuits spontaneously generate varying patterns of neural activity, which profoundly influence sensory responses and behavior¹⁻⁵. These endogenous activity fluctuations are correlated across neural populations and are often quantified by correlations between pairs

of neurons, called noise correlations⁶. Noise correlations are thought to reflect the anatomical circuit connectivity, and they are also dynamically influenced by behavioral and cognitive states^{5,7-9}, in particular, during spatial attention¹⁰⁻¹⁴. Implications of noise correlations for population coding and behavior have been studied extensively¹⁵⁻¹⁹. Yet, how anatomical connectivity and cognitive states interact to define the structure of noise correlations across populations is not well understood.

Spatial selective attention offers a rich experimental domain for studying the combined influence of anatomical connectivity and cognitive factors on the population structure of noise correlations. Changes in noise correlations during attention have been measured across different anatomical dimensions, yielding heterogeneous results. Many studies of noise correlations involved recordings from neurons in different cortical columns, e.g., using rectangular Utah arrays which preferentially sample from laterally separated neurons in more superficial cortical layers^{10,12} (Fig. 1a). These studies found that noise correlations substantially decreased when attention was directed to the receptive fields (RFs) of recorded neurons¹⁰⁻¹². More recent studies used linear multi-electrode arrays to measure attentional modulation of noise correlations within single columns (Fig. 1a) and found effects that varied with layer and area. In V4, noise correlations decreased during attention only in input layers during stimulus-evoked but not spontaneous activity, and no significant changes were observed in superficial and deep layers¹³. In V1, noise correlations decreased only in supragranular layers with no significant changes in granular and infragranular layers¹⁴. In both areas, the magnitude of changes in noise correlations within columns appeared an order of magnitude smaller compared to a sizable reduction of correlations across columns. These data suggest that attentional modulation in correlated variability is not uniform across anatomical dimensions, but depends on lateral distance and cortical layer. The network mechanisms underlying these heterogeneous modulations are unknown.

We hypothesized that heterogeneous changes in noise correlations arise from the modulation of On-Off dynamics propagating through spatially structured cortical connectivity. The On-Off dynamics are spontaneous transitions between phases of vigorous (On) and faint (Off) spiking that occur synchronously across layers of neocortex^{20,21}, and are observed in visual cortex of behaving monkeys^{22,23} (Fig. 1b). The On-Off dynamics reflect the global cortical state associated with arousal and are also modulated locally within retinotopic maps during selective attention^{22,23}. We analyzed spiking activity recorded from single columns in V4 and found that the scale of On-Off dynamics predicted the magnitude of noise correlations and their dependence on lateral distance.

To explain the spatial patterns of noise correlations, we developed a network model of interacting columns with spatially structured connectivity. The key mechanism in our model

is On-Off dynamics that propagate across columns to form spatiotemporal population activity which shapes the structure of noise correlations. Attentional inputs shift the statistics of local On-Off dynamics, which effectively regulates the efficacy of lateral interactions among columns. As a result, attentional inputs modulate the spatial extent of On-Off dynamics, leading to spatially non-uniform changes of noise correlations.

Our model predicts that, during attention, the largest changes in noise correlations occur at intermediate lateral distances (across columns) while within-column changes can be very small. This contrasts with other recently proposed models predicting that attentional changes of noise correlations are spatially uniform²⁴ or monotonically diminish with lateral distance²⁵. To test these competing predictions, we analyzed how changes of noise correlations depend on lateral distance in columnar recordings in V4. Whereas changes in noise correlations were vanishingly small at short distances, they became progressively larger at longer distances, consistent with predictions of our model. Our results provide a unifying framework that explains how heterogeneous patterns of correlated variability emerge within neocortex through interactions of network dynamics and cognitive state.

Results

On-Off dynamics predict the magnitude of noise correlations. We measured spiking activity from all layers within single columns of the visual cortical area V4. Spiking activity was recorded with 16-channel linear array microelectrodes (Fig. 1a) arranged so that receptive fields (RFs) on all channels largely overlapped^{22,26}. During recordings, monkeys performed a spatial attention task, which required detecting changes in the orientation of a visual stimulus in the presence of distractor stimuli. On each trial, an attention cue indicated the stimulus that was most likely to change. In the attention condition, the cue directed animal's attention to the RF stimulus. In the control condition, the cue directed attention to a location outside the RFs of recorded neurons (Methods).

In our columnar recordings, we examined the relationship between the scale of ongoing On-Off dynamics and the magnitude of noise correlations. We quantified the On-Off dynamics by fitting a two-phase Hidden Markov Model (HMM) to the population spiking activity^{22,23} (Fig. 1b, Methods). The HMM models the dynamics of a latent population state that switches between two phases, On and Off, to capture synchronized changes in firing rates across neurons. Spikes on recorded channels were modeled as inhomogeneous Poisson processes with different mean rates during the On and Off phases. The variance explained by a two-phase HMM (R^2)

varied across recording sessions and this variation was tightly correlated with the average noise correlation (Fig. 1c). For most recording sessions (31 total, 67%), the two-phase HMM was the most parsimonious model among HMMs with 1 or up to 8 possible phases²². For the remaining 15 (33%) sessions, a one-phase HMM (i.e. constant firing rates without On-Off transitions) was the most parsimonious model. These one-phase recordings consistently showed lower average noise correlations of multi-unit (MU) activity (mean 0.13) than two-phase recordings (mean 0.32), with a pronounced 59% difference on average (Fig. 1c). Trial-to-trial variability of MU activity, quantified by the Fano factor (FF, the ratio of the spike-count variance to the mean), was also lower in one-phase (mean FF = 1.5) compared to two-phase (mean FF = 2.3) recordings (35% difference). On the other hand, the mean firing rates of MUs were similar between the one-phase (108 Hz) and two-phase (114 Hz) recordings (5% difference). Thus the scale of On-Off dynamics predicted the overall magnitude of correlated variability in our data, which implicates On-Off dynamics as a major source of noise correlations in visual cortex.

Attentional modulation of noise correlations in single columns. We quantified attention-related changes in noise correlations within single columns, separately in superficial (which included granular and supragranular) and deep (infragranular) cortical layers. We combined the granular and supragranular layers because they showed similar changes in noise correlations (Supplementary Fig. 1). In each session, data from each of the recording channels were assigned laminar depth, relative to a common current source density marker²⁶. We found that noise correlations were slightly reduced in superficial and enhanced in deep layers in the attention relative to control conditions (Fig. 1d). To quantify these changes, we calculated a standard modulation index MI_{corr} , which was the difference between noise correlations in the attention and control conditions divided by the sum. In two-phase recordings, the mean $MI_{\text{corr}}^{\text{MU}}$ for MU was -0.029 in superficial layers ($p < 10^{-5}$, Wilcoxon signed rank test, $n = 5, 088$) and 0.022 in deep layers ($p = 0.004$, Wilcoxon signed rank test, $n = 6, 128$) (see Supplementary Tables 1,2 and Supplementary Fig. 2 for a full summary of results). The magnitude and laminar profile (Supplementary Fig. 1) of these noise-correlation changes are consistent with other laminar recordings in V4¹³.

These average changes of noise-correlations within single columns were much smaller than the robust and sizable reduction of noise correlations previously reported for neurons in different columns^{10-12,24}. For comparison, a previous study in V4 using rectangular Utah arrays¹⁰ found that the mean $MI_{\text{corr}}^{\text{MU}}$ was -0.29 , which is an order of magnitude larger than in our columnar recordings. Despite this striking difference in the modulation of noise correlations, the attentional modulation of firing rates and trial-to-trial variability of individual neurons

was similar in our data and the previous study. In two-phase recordings, the mean $MI_{\text{rate}}^{\text{MU}}$ was 0.023 in superficial layers ($p < 10^{-10}$, Wilcoxon signed rank test, $n = 1,752$) and 0.018 in deep layers ($p < 10^{-10}$, Wilcoxon signed rank test, $n = 2,216$), which is more comparable to the previous study¹⁰. Similarly, in two-phase recordings, the mean modulation index of Fano factor $MI_{\text{FF}}^{\text{MU}}$ was -0.010 in superficial layers ($p < 10^{-10}$, Wilcoxon signed rank test, $n = 1,752$) and -0.007 in deep layers ($p < 10^{-10}$, Wilcoxon signed rank test, $n = 2,216$), comparable to the previous study¹⁰. The results were similar in one-phase recordings (Supplementary Table 1, 2). These results suggest that attention-related changes in noise-correlations depend on the relative positions of neurons in the cortex, with sizable changes across columns and minute, layer-dependent changes within columns. Since the striking difference in modulation of noise correlations is not accounted for by differences in the activity of individual neurons, its likely reflects the spatial structure of population dynamics across the cortex.

Network model of interacting cortical columns. We hypothesized that heterogeneous modulations of noise-correlations across layers and columns arise from the On-Off population dynamics and spatial structure of anatomical connectivity in the cortex. To test this hypothesis, we developed a network model of interacting columns with spatially structured connectivity (Fig. 2). The model consists of units interconnected in two parallel two-dimensional lattices, corresponding to the superficial and deep cortical layers (Fig. 2a). Each unit represents a local population of neurons within one layer—superficial or deep—of a single column. Each unit is connected to its four neighboring units in the same layer, mimicking the local structure of horizontal connectivity in the cortex. Visual stimuli and attention are modeled by external inputs to local groups of units.

The key mechanism generating correlated variability in our model is the stochastic On-Off dynamics of the population activity in single columns. In visual cortex of behaving monkeys, the durations of On and Off episodes are distributed exponentially with a timescale of about 100 ms^{22,23}, which indicates that the On and Off phases are metastable with transitions driven by noise²⁷. Accordingly, we model the dynamics of each unit in the network by a two-dimensional dynamical system with two stable fixed points, corresponding to the On and Off phases (Fig. 2b). This dynamical system is a phenomenological mean-field description of a population of excitatory neurons coupled by the vertical recurrent connectivity within the column. The first dynamical variable $r(t)$ represents the mean firing rate of the population. It receives a recurrent self-coupling $F(r)$ and a negative feedback from the second dynamical variable $a(t)$ representing firing-rate adaption^{21,27}. The dynamical system is driven by white noise $\xi(t)$, which causes stochastic transitions between the On and Off fixed points. Each unit also receives

external currents $I_{\text{stim}}(t)$ and $I_{\text{att}}(t)$, which model the bottom-up inputs from visual stimuli and top-down inputs during attention, respectively.

The dynamics of individual units reproduce the On-Off transitions in single columns and their modulation during attention. As in the data from visual cortex^{22,23}, the durations of On and Off episodes in the model are irregular and exponentially distributed (Fig. 2c). In this regime, the dynamics of each unit can be reduced to a two-state Markov process switching between the On and Off phases (Supplementary Note 2.1). The Off-to-On (α_1) and On-to-Off (α_2) transition rates of the Markov process set the average durations of the On and Off episodes: $\tau_{\text{on}} = 1/\alpha_2$ and $\tau_{\text{off}} = 1/\alpha_1$. Consistent with this description, a two-state HMM provides the most parsimonious fit of the On-Off dynamics in our two-phase recordings. Further, our model captures the increase of On-episode durations during attention as observed in the data^{22,23}. During attention, a local group of units representing the attended RFs receives a small excitatory input I_{att} . This attentional input slightly shifts the r -nullcline (Fig. 2b) elevating the threshold for transitioning from the On to Off fixed point, which reduces the On-to-Off transition rate of the Markov process and results in longer average On-episode durations (Fig. 2c).

The horizontal connectivity in the network correlates the On-Off dynamics across units in the lateral dimension. Each unit's firing-rate variable $r(t)$ receives a recurrent excitatory input I_{rec} from its four neighbors on the lattice. As a result, the On-Off dynamics of each unit are influenced by the activity of its neighbors. The more neighbors in the On phase, the larger is the excitatory input I_{rec} , which elevates the threshold for On-to-Off and lowers the threshold for Off-to-On transitions. In the description of a two-state Markov process, this is equivalent to a dependence of the Off-to-On and On-to-Off transition rates on the On/Off phases of the neighbors: $\alpha_1 + \beta S_{\pm}$ and $\alpha_2 - \beta S_{\pm}$ (Supplementary Note 2.3). Here, the variable S_{\pm} indicates the number of neighbors in the On phase at each time, and β is the effective coupling strength that depends on the parameters of the dynamical system as well as on external inputs (Supplementary Note 2.3). The reduced network of coupled binary On-Off units follows Glauber dynamics²⁸, allowing us to calculate noise correlations analytically in our model. In simulations, both the dynamical-system and binary-unit versions of the network exhibit similar spatiotemporal dynamics, where the On/Off phases form local spatial clusters (Fig. 2d).

We model two sources of spiking variability: the On-Off fluctuations of the population activity and stochasticity of spike generation in individual neurons^{21,29}. We simulate spikes of individual neurons as inhomogeneous Poisson processes with different mean rates during the On and Off phases generated by the network (Fig. 2e). This doubly-stochastic description coincides with the assumptions of the HMM used to fit the experimental data. We match the model parameters to the experimental data by fitting the data with the HMM, which provides us

the estimates of the On-Off transition rates (α_1 and α_2) and the On and Off firing rates (r_{on} and r_{off}) for all MUs and single units (SUs) in each recording session and task condition (Fig. 2f). We then use these parameters to predict noise correlations and compare these predictions with the values measured for the same neuron pairs in the data.

The model accounts for correlated variability in single columns. We used the two-phase recordings to test how accurately the model predicts changes of correlated variability in single columns during attention. In the model, the On-Off dynamics are the source of correlations between responses of individual neurons. All neurons in the population represented by a single unit (column) follow the same shared sequence of On and Off episodes. The spiking responses differ across neurons because of the independent Poisson noises as well as differences in their On and Off firing rates. We derived analytical formulas for the Fano factor and noise correlations, measured over an arbitrary time-window T , as functions of the model parameters: the On-Off transition rates and the On and Off firing rates of each neuron (Methods and Supplementary Note 2.1). The model analytically predicts the dependence of Fano factor and noise correlations on the measurement time-window⁶, indicating that this dependence is determined by the timescales of On-Off dynamics (Supplementary Note 2.1). We used these analytical formulas with the parameter values estimated from the data by the HMM to predict the FF and noise correlations for, respectively, each neuron and neuron pair in our dataset. We compared these model predictions with direct measurements from the experimental data.

Our model makes a specific prediction that the key factor determining the magnitude of FF and noise correlations in single columns is the On-Off firing-rate difference $\Delta r = r_{\text{on}} - r_{\text{off}}$. Specifically, FF is directly proportional to Δr , and the noise correlation between neurons i and j is proportional to the product $\Delta r_i \Delta r_j$ (Methods and Supplementary Note 2.1). This dependence on Δr is intuitive, because the source of correlations within a column is the shared On-Off switching, hence the stronger a neuron is modulated by the On-Off dynamics (the greater is Δr), the stronger it will be correlated with other neurons in the same column. The dependence of FF and noise correlations on Δr is evident in an example recording (Fig. 3a,b), where different MUs exhibit a variety of On-Off firing-rate differences Δr . The FF ranges broadly across MUs, from ~ 1 up to ~ 9 , and this variation is very well predicted by Δr (Fig. 3a). Similarly, MU pairs with the largest product $\Delta r_i \Delta r_j$ also exhibit the largest noise correlations (Fig. 3b). While both FF and noise correlations also weakly depend on the On-Off transition rates, the On-Off firing-rate difference is the main factor defining the broad distributions of these quantities in single columns (Supplementary Fig. 3).

As a consequence, the model also predicts that the changes in FF and noise correlations

during attention are proportional to the changes in Δr and $\Delta r_i \Delta r_j$, respectively. This prediction was clearly borne out by the data: the measured change in FF had a strong trend as a function of the change in Δr (y-axis vs. color-axis in Fig. 3c), and the measured change in noise correlations had a strong trend as a function of the change in $\Delta r_i \Delta r_j$ (y-axis vs. color-axis in Fig. 3d). Moreover, changes in the FF and noise correlations measured from the data were accurately matched by the model predictions (y- vs. x-axis in Fig. 3c,d). The attention-related changes of noise correlations range widely across the population, with noise correlations substantially reduced and enhanced in many pairs. This entire broad distribution is accurately matched by the analytical predictions of the On-Off dynamics model in single columns. Despite substantial changes in many pairs, the average change of noise correlations in single columns is near zero (Fig. 1b), since the changes of $\Delta r_i \Delta r_j$ are broadly distributed but close to zero on average. Thus, our model of On-Off dynamics explains the observed changes in correlated variability during attention in single columns.

Decay of noise correlations with lateral distance. Next, we analyzed the dependence of noise correlations on the lateral distance in our laminar recordings and in the network model. Previous studies in the visual cortex found that noise correlations decrease with the lateral distance^{24,30,31}. These studies used multi-electrode arrays with lateral spacing between electrodes ranging from ~ 0.35 to 4 mm, i.e. sampling distant neuron pairs in different columns. With the laminar recordings, we tested how noise correlations depend on the lateral distance over a much shorter range of distances within single or nearby columns. We leveraged the fact that laminar recordings generally exhibit slight horizontal shifts due to variability in the penetration angle (Fig. 4a). As a surrogate for horizontal displacements between pairs of channels, we used distances between centers of their RFs. To estimate the range of physical distances in the cortex spanned by our laminar recordings, we converted the RF-center distances to cortical distances using the cortical magnification factor for each eccentricity³². The range of distances spanned by our recordings was $\sim 4 - 6$ dva or ~ 1.5 mm (Methods and Supplementary Fig. 4).

We found that noise correlations decreased with lateral distance in two-phase but not in one-phase recordings. In two-phase recordings, noise correlations monotonically decreased with the RF-center distance both in superficial and deep layers (Fig. 4b, linear regression, one-sided t-test, slope -0.09 ± 0.01 , $p < 10^{-8}$ in superficial layers, slope -0.04 ± 0.01 , $p < 10^{-3}$ in deep layers). With the conversion to cortical distances³², noise correlations also decreased continuously with the lateral cortical distance over $\lesssim 1$ mm range (Supplementary Fig. 4). Note that most pairs in our data were at very short distances, with the median estimated cortical distance 0.72 mm. Thus the decay of noise correlations with lateral distance spans all distances

within nearby and across distant columns. In contrast, noise correlations did not decrease with the RF-center distance in the one-phase recordings (Fig. 4c, linear regression, one-sided t-test, slope -0.010 ± 0.008 , $p = 0.1$ in superficial layers, slope 0.014 ± 0.004 , $p = 0.997$ in deep layers). These results suggest that On-Off dynamics give rise to the lateral distance-dependence of noise correlations.

In our network model, the dependence of noise correlations on the lateral distance arises from the spatiotemporal On-Off dynamics. Whereas all neurons represented by a single unit in the network follow the same shared sequence of On-Off phases, neurons represented by different units follow their respective On-Off sequences. The On-Off dynamics are correlated across nearby units in the lateral dimension due to horizontal connectivity in the network. Since the horizontal connections are spatially local and relatively weak, the synchrony of On-Off dynamics is not global across the entire network, but localized to a finite range of lateral distances. Thus the On/Off phases form spatial clusters with a characteristic spatial length scale, and beyond this spatial scale the On-Off phases are uncorrelated. This network mechanism leads to a continuous decrease of noise correlations with the lateral distance in the model (Fig. 4d).

Using the binary-unit reduced network model, we derived an analytical formula for the dependence of noise correlations on the lateral distance d (Methods). Our calculations show that noise correlations decay with the distance exponentially as $\mathcal{A} \exp(-d/L)$. This formula describes noise correlations both within and across columns. Within columns ($d = 0$), the formula reduces to the pre-factor $\mathcal{A} = \mathcal{A}(\alpha_1, \alpha_2, \Delta r_i, \Delta r_j)$, which accounts for the dependence of noise correlations on the On-Off transition rates and the On-Off firing-rate difference, as described in the previous section. Across columns (finite $d > 0$), the formula accounts for the spatial structure of noise correlations with the exponential discount factor $\exp(-d/L)$. The space-constant L of this exponential decay, termed correlation length, depends on the On-Off transition rates and on the effective coupling strengths β between units in the network: $L = \sqrt{\beta/(\alpha_1 + \alpha_2)}$ (in dimensionless units of the lattice constant, see Methods and Supplementary Note 2.3). This analytical result agrees well with simulations of the full dynamical-system network model (Fig. 4d).

The exponential decay of noise correlations with distance in our model, characterized by the correlation length L , is consistent with the decrease in noise correlations observed over a wide range of cortical distances spanned by our laminar (Fig. 4b) and previous lateral recordings^{24,30,31} from the primate visual cortex. Our model can also reconcile the decay of noise correlations with distance in lateral recordings^{24,30,31} with the lack of distance dependence in the one-phase recordings. With some heterogeneity, if a random fraction of units in the model does not exhibit On-Off transitions (due to a more stable fixed point), the activity of

these one-phase units is not correlated with other units at all distances. Thus lateral sampling from a mixture of one-phase and two-phase phase units would uniformly lower the average of noise correlations without affecting their distance dependence.

Differential changes in noise correlations arise from attentional modulation of the correlation length. In our network model, attentional inputs restructure the spatiotemporal On-Off dynamics, leading to differential changes of noise correlations within versus across columns. In the network simulations, a local group of units receives an attentional input I_{att} , while other units without this input ($I_{\text{att}} = 0$) are in the unattended control condition (Fig. 2d). With an excitatory attentional input ($I_{\text{att}} > 0$), average noise correlations between neurons in the same column change very little relative to control ($\text{MI}_{\text{corr}} = -0.03$), while noise correlations between neurons in different columns are substantially reduced ($\text{MI}_{\text{corr}} = -0.28$, Fig. 5a). These results replicate the order-of-magnitude difference of MI_{corr} observed between laminar versus lateral recordings from the visual cortex. We repeated simulations for a range of excitatory ($I_{\text{att}} > 0$) and inhibitory ($I_{\text{att}} < 0$) attentional inputs. The excitatory inputs reduced noise correlations, whereas inhibitory inputs increased noise correlations, but in all cases the average changes of noise correlations within columns were very small compared to sizable changes across columns of the network (Fig. 5b).

To reveal the mechanism leading to differential changes of noise-correlations within versus across columns, we examined how attentional inputs affect the dependence of noise correlations on the lateral distance in our network. In simulations, excitatory attentional inputs produce a faster decay of noise correlations with lateral distance, which corresponds to a shorter correlation length ($L_{\text{att}} < L_{\text{ctl}}$, Fig. 5c). Due to this faster spatial decay, noise correlations at intermediate lateral distances (finite $d > 0$) are considerably lower in attention relative to control condition, even though changes of noise-correlations within columns ($d = 0$) are small. Inhibitory attentional inputs, on the other hand, produce a slower decay of noise correlations with lateral distance, which corresponds to a longer correlation length ($L_{\text{att}} > L_{\text{ctl}}$) and results in higher noise correlations at intermediate lateral distances. Thus changes of the correlation length L produce sizable changes of noise correlations at intermediate lateral distances (across columns) even when noise correlations within columns do not change.

To understand the network mechanism by which attentional inputs modulate the correlation length, we leveraged the analytical formula $L = \sqrt{\beta/(\alpha_1 + \alpha_2)}$. In the dynamical-system model, an excitatory attentional input shifts the r -nullcline, which increases Off-to-On (α_1) and decreases On-to-Off (α_2) transition rates (Fig. 2b,c). Since α_1 and α_2 change only moderately and in opposite directions, their sum remains nearly constant (Fig. 5d). Therefore changes of the

correlation length L are mainly driven by changes in the effective coupling strength β , which decreases steeply with an increasing attentional input (Fig. 5e). The effective coupling strength β decreases because an excitatory attentional input stabilizes the On fixed point, thereby effectively reducing the efficacy of the lateral recurrent inputs to drive the On-Off transitions. Vice versa, an inhibitory attentional input makes the On fixed point less stable, thereby enhancing the relative efficacy of the lateral recurrent inputs and hence extending the spatial correlation length in the network (Supplementary Note 3.2). Thus the attentional input modulates the correlation length by regulating the relative efficacy of lateral interactions between columns³³, which leads to differential changes of noise correlations within versus across columns.

The model predicts distance-dependent changes of noise correlations. The major changes of noise correlations in our model are driven by changes of the correlation length L . The model therefore makes a specific prediction that changes of noise correlations during attention are not uniform across space. Noise correlations decay exponentially with the lateral distance, with different decay rates in attention and control conditions. Hence the spatial profile of noise-correlation changes is defined by the difference of two exponential decays: $\exp(-d/L_{\text{att}})$ and $\exp(-d/L_{\text{ctl}})$. At very short lateral distances (within columns), average changes of noise-correlations are minute (Fig. 1d). At very long lateral distances, the average changes are negligible, because the overall magnitude of noise correlations vanishes. The largest changes of noise correlations are predicted to occur at intermediate lateral distances, where the difference between two exponential decays dominates. Thus the network mechanism in our model predicts a non-monotonic dependence of noise-correlation changes on lateral distance (Fig. 6a). This prediction contrasts with predictions of recently proposed alternative models, which instead predict that attention-related changes of noise correlations are spatially uniform²⁴ or monotonically diminish with lateral distance²⁵ (Supplementary Note 3). Therefore examining the spatial profile of noise-correlation changes in the data could distinguish among network mechanisms these alternative models postulate.

We analyzed how changes of noise correlations during attention depend on the lateral distance (estimated by the RF-center distance) in our laminar recordings. As predicted by our model, changes of noise correlations in the two-phase recordings were not uniform across space. Changes of noise correlations were smallest at very short lateral distances and became progressively greater at longer distances, with opposite trends in the superficial and deep layers (Fig. 6b). In the superficial layers, noise correlations decreased during attention, with greater reduction at longer distances (linear regression, one-sided t-test, slope -0.017 ± 0.004 , $p < 10^{-3}$). The extrapolation of this trend to intermediate lateral distances ($1 \leq d \leq 4$ mm) is consistent

with a robust reduction of noise correlations during attention observed in Utah-array recordings, which also sample from superficial layers^{10,12}. In the deep layers, the attentional effects on noise correlations were reversed from that in superficial layers (linear regression, one-sided t-test, slope 0.006 ± 0.004 , $p = 0.06$; t-test for slopes superficial versus deep: $p = 0.6 \cdot 10^{-3}$). In contrast, changes of noise correlations did not depend on lateral distance in the one-phase recordings (Fig. 6c, linear regression, t-test, slope -0.0005 ± 0.0061 , $p = 0.93$ in superficial layers and slope -0.01 ± 0.01 , $p = 0.35$ in deep layers). These results indicate that On-Off dynamics give rise to distance-dependence of noise-correlation changes during attention. The spatial profile of noise-correlation changes in two-phase recordings is consistent with our network mechanism, but inconsistent with the previously proposed models, which predict either spatially uniform²⁴ or monotonically diminishing changes with greatest changes at zero distance²⁵. The observed spatial profiles of noise-correlation changes indicate that the correlation length decreases in superficial and increases in deep layers, which suggests that superficial and deep layers receive different modulatory inputs during attention³⁴.

Discussion

Our results show that On-Off dynamics are a major source of correlated variability in the visual cortex. We found that On-Off dynamics predicted the magnitude of noise correlations, their attentional modulation and dependence on lateral distance. The average changes of noise correlations during attention were very small within single columns of the visual area V4. Noise correlations slightly decreased in superficial and increased in deep layers, but the changes were an order of magnitude smaller than a robust and sizable reduction of noise correlations between neurons in different columns. Since noise correlations strongly impact the amount of sensory information encoded by the population^{15–19}, a reduction of noise correlations is thought to be a major contributor to the improved psychophysical performance during attention¹⁰. Our results show, however, that changes of noise correlations are not uniform: their magnitude and sign depend on the relative anatomical positions of neurons within layers and columns of the visual cortex. These heterogeneous changes of noise correlations may reflect unique contributions to behavioral improvements from different functional groups of neurons defined by their anatomical positions within the circuit.

To explain differences in attention-related changes of noise correlations within versus across columns, we developed a network model of interacting cortical columns. The key mechanism generating correlated variability in the model is On-Off dynamics, metastable transi-

tions between the high and low firing-rate fixed points in single columns. Spatially structured horizontal connectivity in the network correlates On-Off dynamics across nearby columns to produce an exponential decay of noise correlations with lateral distance, characterized by the correlation length. Attentional inputs restructure the spatiotemporal On-Off dynamics and modulate the correlation length, which results in distance-dependent changes of noise correlations. The model accurately captures attention-related changes of noise correlations in our data. Moreover, it makes a testable prediction that greatest changes of noise correlations occur at intermediate lateral distances. Consistent with this prediction, we found that the magnitude of noise-correlation changes gradually increased with lateral distance in our laminar recordings. This experimental finding is not consistent with other models which predict that changes of noise correlations are spatially uniform²⁴ or monotonically diminish with lateral distance²⁵.

Mechanisms generating correlated variability in recurrent networks. Several mechanisms were proposed to explain how correlated fluctuations arise in cortical networks. There are two general classes of models: one relies on external shared variability and another generates shared variability via intrinsic network dynamics. In models with external shared variability, the source of correlated fluctuations is assumed to be outside the network, and the network merely filters the correlated input noise^{25,35,36}. In most of these models, the mechanism is based on a spatial connectivity structure that locally breaks the excitation-inhibition balance. The classical balanced network model with random connectivity^{37,38} operates in an asynchronous regime, where the tight excitation-inhibition balance cancels any input correlations resulting in zero average noise correlations³⁹. The spatial connectivity structure, where recurrent inhibition is broader than feedforward excitation, breaks the balance locally, hence the input correlations cannot be canceled resulting in positive average noise correlations^{35,36}. However, to match the experimentally observed temporal and spatial scales of correlations, all of these models have to assume *ad hoc* spatiotemporal structure of the input noise^{25,35,36}.

The second class of models can generate shared variability internally. One mechanism is based on breaking stability in some spatial Fourier modes in a spatially organized network. For example, in a two-dimensional balanced network with slow inhibitory kinetics, shared fluctuations arise from instability at some spatial frequency that generates rate chaos²⁴. Similarly, in a one-dimensional balanced ring model, strong correlations arise from a feed-forward structure in some Fourier modes of connectivity⁴⁰. However, in these models correlations arise from fluctuations around a global fixed point with a timescale defined by the mismatch between excitatory and inhibitory synaptic time-constants, i.e. just a few milliseconds. This fast timescale is inconsistent with data as fluctuations of cortical activity occur on a timescale of about hundred

milliseconds^{22,41}.

An alternative mechanism that can account for the slow timescale of cortical fluctuations is based on multi-stability. In this case, slow correlated fluctuations arise from stochastic transitions between multiple fixed points in the network. Multi-stability can result from clustered excitatory connectivity, where each cluster corresponds to a fixed-point attractor⁴². Further, bistability between high and low firing-rate attractors can arise in unstructured networks with strong recurrent excitation and slower negative feedback such as firing-rate adaptation^{21,27,43,44} or short-term synaptic depression⁴⁵. The models with multi-stability capture the slow timescale of cortical fluctuations and produce realistic noise correlations^{21,27,44}. However, the multi-stable networks studied previously were not endowed with a spatial connectivity layout akin to organization of cortical networks, hence they do not produce any spatial structure of noise correlations.

To account for both the slow timescale and spatial structure of noise correlations in the visual cortex, our network model combines local bistable On-Off dynamics with spatially organized connectivity. In our model, correlated variability arises from metastable transitions between the On and Off fixed points, and not from fluctuations around a single global attractor²⁵. Our results suggest that a theory of noise correlations in the visual cortex should incorporate both the anatomical connectivity structure of visual areas as well as local bistability of population dynamics in single columns.

Mechanisms of attentional modulation of noise correlations. Recurrent network models were previously developed to suggest possible circuit mechanisms that produce a reduction of noise correlations during attention^{24,25,46}. These models are based on a dynamical mechanism, where the network operates around a global fixed point and attentional inputs increase the stability of this fixed point leading to suppression of correlated fluctuations. Specifically, in the network with intrinsically generated shared variability, the stability of the operating point can be increased by up-regulating activities of inhibitory neurons²⁴. However, elevated inhibition reduces firing rates of excitatory neurons, which contradicts attentional enhancement of firing rates in experiments. In the network filtering external noise, the stability of the global fixed point can be increased by excitatory inputs when the network operates in inhibition dominated regime²⁵. In this scenario, an excitatory input increases effective lateral connectivity, which suppresses the transmission of the correlated input noise (Supplementary Note 3.4).

The mechanism we propose differs from these previous models. First, we show that a reduction of noise correlations during attention is not universal. Therefore, a network mechanism should account for heterogeneous changes of noise correlations across different anatomical

ical dimensions. Second, the mechanism in our model is based on local bistability of On-Off dynamics in single columns. Attentional inputs change the stability of the On and Off fixed points, which effectively modulates the efficacy of lateral interactions across the network leading to changes of the correlation length. This mechanism is fundamentally distance-dependent, as the major changes of noise correlations in our model are driven by changes of the correlation length (Supplementary Note 2.3, 3.2). As a consequence, we find that in Fourier space the lower spatial frequency modes contribute most to noise-correlation changes (Supplementary Note 3.2). This result partially agrees with the previous model²⁴, where the dominant part of noise-correlation changes arises from zero spatial frequency mode, which however, predicts a spatially uniform modulation of noise correlations. In contrast, contributions from higher spatial frequency modes are not negligible in our model. A combination of all spatial frequency modes generates a non-monotonic profile of noise-correlation changes in lateral dimension, a prediction that was confirmed in our data.

Several biophysical substrates could mediate the network mechanism of attentional modulation in our model. Top-down projections from frontal cortical areas^{47,48} can provide temporally and spatially precise inputs to drive fast and local modulation of On-Off dynamics in the visual cortex. Neuromodulatory inputs can also mediate effects of attention⁴⁹ and can influence On-Off dynamics by modulating neural excitability and firing-rate adaptation⁵⁰. Further, laminar distribution of top-down inputs³⁴ and of neuromodulation, combined with layer-specific horizontal connectivity could account for the differential modulation of noise correlations in superficial and deep layers that we observed. Identifying precise mechanisms by which these multiple biophysical components interact within a columnar microcircuit is an important direction for future work.

Methods

Behavior and electrophysiology. Experimental procedures have been described previously^{22,51}. Two male monkeys (*Macaca mulatta*, 8 – 12 kg) were used in experiments. Experimental procedures were in accordance with NIH Guide for the Care and Use of Laboratory Animals, the Society for Neuroscience Guidelines and Policies, and Stanford University Animal Care and Use Committee.

In the attention task, the monkey reported cued orientation changes of visual stimuli with an antisaccade response. On each trial, a small central cue indicated the stimulus that was most

likely to change. The cued stimulus was therefore a target of covert attention. The monkeys reported stimulus changes with an antisaccade to the location opposite to the change, which was therefore a target of overt attention due to anticipation of antisaccadic response^{22,51}. Modulations of neural responses in V4 were highly similar during the covert and overt attention, including changes in firing rates, spiking variability and noise correlations^{22,51}, and therefore we combined the covert and overt attention conditions in our analyses. The monkey initiated each trial by fixating a central fixation dot on the screen. Within several hundred milliseconds, four peripheral stimuli appeared (static Gabor patches: oriented black and white gratings in a circular Gaussian aperture). After a short delay, the attention cue appeared: a short line originating at the fixation dot and extending in the direction of one of the four stimuli, randomly chosen on each trial with equal probability. The cue indicated with $\sim 90\%$ validity which of the four stimuli, if any, would change on each trial. After a post-cue period of 600 – 2,300 ms, all stimuli synchronously disappeared for a brief interval and then reappeared. On half of the trials, one of the four stimuli reappeared with a changed orientation (i.e. rotated in place), and the monkey was rewarded for performing a saccadic eye movement to the location opposite to the changed stimulus. On the other half of the trials, all stimuli reappeared with the same orientations as they had before disappearing, and the monkey was rewarded for maintaining fixation on the central dot.

While monkeys performed the attention task, recordings were made in the visual area V4 with a 16-channel linear array microelectrodes^{22,51}. The total length of array is 2.25 mm, and the spacing between electrical contacts is 150 μm . Recordings were targeted with MRI to be as perpendicular to cortical layers as possible so as to maximize the overlap of receptive fields (RFs) of recorded neurons. Each of the recording channels were assigned laminar depth relative to a common current source density marker as described previously²⁶.

Data analysis. We measured Fano factor and noise correlations in our recordings using spike-counts N of MUA and SUA in 200 ms bins (400 to 600 ms window after the attention cue onset). The Fano factor is the ratio of the spike-count variance to its mean across trials: $\text{Var}[N]/\text{E}[N]$. The noise correlation r_{sc} is the Pearson correlation coefficient between spike-counts N_i and N_j of two neurons:

$$r_{\text{sc}} = \frac{\text{E}[N_i N_j] - \text{E}[N_i]\text{E}[N_j]}{\sqrt{\text{Var}[N_i]\text{Var}[N_j]}}. \quad (1)$$

We estimated parameters of the On-Off dynamics in single columns by fitting population spiking activity in our recordings with a two-state Hidden Markov Model (HMM) as described previously²². HMM has a latent variable representing an unobserved population state

that stochastically switches between the On and Off phases following Markov dynamics. Spikes on 16 simultaneously recorded channels are assumed to be generated by inhomogeneous Poisson processes, with different mean rates during the On and Off phases. The latent On-Off state is shared by the population, but the On and Off firing rates can differ across neurons. HMM was fitted to MUA spike-counts in 10 ms bins, during a time-window starting at 400 ms after the attention cue onset and until the end of the post-cue period. The duration of this time-window ranges between 200 – 1,900 ms across trials. HMM was fitted separately for each of the 32 task conditions (4 attention conditions \times 8 grating orientations). The HMM parameters were optimized with the Expectation-Maximization algorithm²². The HMM had 34 parameters: firing rates in the On (r_{on}) and Off (r_{off}) phases for each of 16 channels and transition probabilities p_{on} and p_{off} for the entire population.

We estimated lateral shifts between channels in our laminar recordings by distances between centers of their RFs. The RF mapping procedure was described previously²². RFs were measured by recording spiking responses to briefly flashed stimuli on an evenly spaced 6×6 grid covering the lower left visual field. Spikes in the window 0 to 200 ms relative to stimulus onset were averaged across all presentations of each stimulus. The RF center was defined as the center of mass of the response map. The lateral cortical distance d_{cortical} (mm) was estimated from the RF-center distance d_{RF} (d.v.a) using the cortical magnification factor M for each eccentricity E ³²:

$$d_{\text{cortical}}(\text{mm}) = 9 - M \times d_{\text{RF}}(\text{d.v.a}), \quad M = 3.1E^{-0.9}. \quad (2)$$

Network model of interacting columns. The model describes spatiotemporal dynamics of neural population activities across cortical surface. The network consists of two two-dimensional square lattices of units, representing superficial and deep cortical layers. Each unit represents a local population of neurons within one layer of a single column. The dynamical variable $r(\mathbf{x}, t)$ represents the mean firing-rate of this population. The two-dimensional lateral coordinates are denoted as \mathbf{x} . The dynamics of the network model are given by

$$\begin{aligned} \epsilon \frac{d}{dt} r &= F(r) - a + W \nabla^2 r + I_{\text{stim}} + I_{\text{attn}}, \\ \frac{d}{dt} a &= gr - a + f + \sqrt{2Q} \xi. \end{aligned} \quad (3)$$

Here $a(\mathbf{x}, t)$ is the adaption variable, ξ is a white Gaussian noise of unity intensity, and we omit

the spatial indexes of variables r and a for clarity. The function $F(r)$ is given by

$$F(r) = \begin{cases} -1 - r & , & r \leq -1/2 \\ r & , & -1/2 < r < 1/2 \\ 1 - r & , & r \geq 1/2 \end{cases} . \quad (4)$$

This piecewise linear function approximates the inverted N-shaped r -nullcline, typically used in rate-models with adaptation²¹, which allows us to analytically reduce the dynamical system to a binary-unit model⁵². The term $W\nabla r$ represents lateral interactions between neighboring units, where $\nabla^2 r = \partial_x^2 r + \partial_y^2 r + 2\partial_x \partial_y r$ implements a diffusive coupling and W is the interaction strength parameter. The external currents I_{stim} and I_{att} are applied to local groups of units to model stimulus and attentional inputs, respectively. A constant $\epsilon \ll 1$ separates the timescales of the fast firing-rate variable r and slow adaption variable a . The parameters g, f, Q are chosen so that the system is bi-stable⁵², where the population rate r stochastically switches between two stable fixed points, corresponding to the On and Off phases.

We match the model to experimental data using the fitted HMM parameters. Specifically, the HMM transition matrix P ($p_{11} = p_{\text{off}}, p_{12} = 1 - p_{\text{off}}, p_{22} = p_{\text{on}}, p_{21} = 1 - p_{\text{on}}$) provides an estimate of the On-Off transition rates: $\alpha_1 = (1 - p_{\text{off}})/\Delta t$ and $\alpha_2 = (1 - p_{\text{on}})/\Delta t$, where $\Delta t = 10\text{ms}$ is the bin size used for HMM fitting. HMM also estimates the On and Off firing rates r_{on} and r_{off} for each MUA and SUA, which we use to generate spikes of the model neurons. To this end, for each network unit we segment the simulated time-series $r(t)$ into the On ($S = 1$) and Off ($S = 0$) phases as $S(t) = \Theta[r(t)]$, where Θ is the Heaviside step function. The spike counts are then generated from inhomogeneous Poisson processes with rates $r_i(t)$, where the firing rate for neuron i is

$$r_i(t) = r_{\text{off},i} + \Delta r_i S_i(t) , \quad \Delta r_i = (r_{\text{on},i} - r_{\text{off},i}) . \quad (5)$$

Simulations. We simulated the network model Eq. (3) on a 256×256 discrete square lattice with a time step of 0.01 s. The unit activities are initialized randomly. We compute noise correlations from 100 simulated trials for each set of parameters. On each trial, we simulated the period of spontaneous activity, stimulus period and attention-cue period, as in the experimental data. During stimulus period, external inputs I_{stim} were applied to two local groups of units with the size 50×50 . During the attention-cue period, one of these two groups also received attentional inputs I_{att} . To calculate noise correlations, we either assigned fixed values of r_{on} and r_{off} or sampled them from distributions of r_{on} and r_{off} extracted from experimental data by HMM.

Reduction to a binary-unit network. When the dynamical-system network operates in the bistable regime, the activity of each unit i can be approximated by a binary variable S_i ⁵², where $S_i = 1$ refers to On phase, and $S_i = 0$ to Off phase. We derived a reduced network model, where the dynamical equations describe the state transition probabilities of binary units. Using the mean field approximation, we derived an approximate form for transition rates of binary units (Supplementary Note 2.3). In the leading approximation order, we have

$$w(S_i = 0 \rightarrow 1) \approx \alpha_1 + \beta_1 S_{i\pm 1}, \quad w(S_i = 1 \rightarrow 0) \approx \alpha_2 - \beta_1 S_{i\pm 1}. \quad (6)$$

Here $S_{i\pm 1}$ are the sum of activities of neighboring units that are connected to a given unit S_i . α_1 , α_2 and β_1 are functions of parameters in the dynamical-system model: f , g , Q , I_{stim} and I_{attn} . This reduced model allows us to derive analytical formulas for correlations between units in the network.

The reduced network model of binary units. The binary-unit network operates on a two-dimensional square lattice. The network consists of N units. Each unit can be in a discrete On ($S_i = 1$) or Off ($S_i = 0$) state, represented by a binary variable $S_i = \{0, 1\}$, ($i = 1, \dots, N$). At time t , the probability of the system to be in a certain configuration $\{S\} = \{S_1, S_2, \dots, S_N\}$ is denoted as $P(\{S\}, t)$. The rate of change of $P(\{S\}, t)$ is described by the master equation:

$$\frac{d}{dt}P(\{S\}, t) = -P(\{S\}, t) \sum_i w(S_i) + \sum_i P(\{S\}^{i*}, t)w(1 - S_i). \quad (7)$$

Here $\{S\}^{i*} = \{S_1, S_2, \dots, 1 - S_i, \dots, S_N\}$, and $w(S_i)$ is the transition rate. When $S_i = 0$, the transition rate of S_i from 0 to 1 is

$$w(S_i = 0) = \alpha_1 + \beta_1(S_{i\pm 1}). \quad (8)$$

When $S_i = 1$, the transition rate of S_i from 1 to 0 is

$$w(S_i = 1) = \alpha_2 - \beta_2(S_{i\pm 1}). \quad (9)$$

Here α_1 and α_2 represent the baseline transition rates of each unit without interactions with other units, and $\beta_{1,2}$ describe how the transition rates are influenced by nearby units $S_{i\pm 1}$. The diffusive coupling between units is described by the discrete Laplacian:

$$S_{i\pm 1} = S_{i+1} - S_i + S_{i-1} - S_i. \quad (10)$$

For simplicity, we use a single index i to represent indices in arbitrary dimension. For example, in two dimensions $i = (x, y)$, and we have

$$S_{i\pm 1} = S_{x+1,y} - S_{x,y} + S_{x-1,y} - S_{x,y} + S_{x,y+1} - S_{x,y} + S_{x,y-1} - S_{x,y}. \quad (11)$$

Based on the master equation, the dynamics of the first and second moments are given by

$$\begin{aligned}\frac{d}{dt}\langle S_i \rangle(t) &= \alpha_1 - (\alpha_1 + \alpha_2)\langle S_i \rangle + \beta_1\langle S_{i\pm 1} \rangle, \\ \frac{d}{dt}\langle S_i S_j \rangle(t) &= \alpha_1(\langle S_i \rangle + \langle S_j \rangle) - 2(\alpha_1 + \alpha_2)\langle S_i S_j \rangle + \beta_1(\langle S_{i\pm 1} S_j \rangle + \langle S_{j\pm 1} S_i \rangle).\end{aligned}\quad (12)$$

We studied the dynamics of the binary-unit network analytically and in simulations. In simulations, the states of all units were updated based on their transition rates in 10 ms time bins.

Theoretical prediction of noise correlations. Assuming the network evolved to the equilibrium state, we derived in the continuum limit the steady-state solution for the averaged first moment $S(\infty)$ and quadratic moments $G(d; \infty)$:

$$S(\infty) = \frac{\alpha_1}{\alpha_1 + \alpha_2}, \quad (13)$$

$$G(d; \infty) = [S(\infty)]^2 + S(\infty)(1 - S(\infty)) \exp\left(-\frac{d}{L}\right). \quad (14)$$

Here the dimensionless correlation length L is given by

$$L = \sqrt{\frac{\beta_1}{\alpha_1 + \alpha_2}}, \quad (15)$$

and d is the dimensionless lateral distance measured in units of the lattice constant Δd .

Using these expressions for the first moment $S(\infty)$ and quadratic moments $G(d; \infty)$, we derived an analytical formula for the noise correlations. Consider a pair of neurons (\mathbf{x}, i) and (\mathbf{y}, j) that are indexed by the lateral positions \mathbf{x}, \mathbf{y} of units to which they belong, and by their indices i, j within these units. Spike-counts $N(\mathbf{x}, i)$ and $N(\mathbf{y}, j)$ of these two neurons are measured in a time-window of duration T . The theoretical prediction of noise correlation $r_{sc}[N(\mathbf{x}, i), N(\mathbf{y}, j)]$ is given by

$$r_{sc}[N(\mathbf{x}, i), N(\mathbf{y}, j)] = \mathcal{A}(\alpha_1, \alpha_2) \exp\left(-\frac{|\mathbf{x} - \mathbf{y}|}{L}\right). \quad (16)$$

This equation shows that noise correlations decay exponentially with the lateral distance $d = |\mathbf{x} - \mathbf{y}|$, with the decay-rate characterized by the correlation length L . The amplitude $\mathcal{A}(\alpha_1, \alpha_2)$ depends on the On-Off transition rates α_1, α_2 , and on the On/Off firing rates $r_{\text{off}}(\mathbf{x}, i), r_{\text{off}}(\mathbf{y}, j), \Delta r(\mathbf{x}, i), \Delta r(\mathbf{y}, j)$. Specifically,

$$\mathcal{A}(\alpha_1, \alpha_2) = \frac{V(\alpha_1, \alpha_2)\Delta r(\mathbf{x}, i)\Delta r(\mathbf{y}, j)}{\sqrt{\text{Var}[N(\mathbf{x}, i)]\text{Var}[N(\mathbf{y}, j)]}}, \quad (17)$$

$$\begin{aligned}\text{Var}[N(\mathbf{x}, i)] &= (\Delta r(\mathbf{x}, i))^2 V(\alpha_1, \alpha_2) + r_{\text{off}}(\mathbf{x}, i)T + \frac{\alpha_1}{\alpha_1 + \alpha_2} T \Delta r(\mathbf{x}, i), \\ \text{Var}[N(\mathbf{y}, j)] &= (\Delta r(\mathbf{y}, j))^2 V(\alpha_1, \alpha_2) + r_{\text{off}}(\mathbf{y}, j)T + \frac{\alpha_1}{\alpha_1 + \alpha_2} T \Delta r(\mathbf{y}, j),\end{aligned}\quad (18)$$

where

$$V(\alpha_1, \alpha_2) = \frac{2(\alpha_1 \alpha_2)}{(\alpha_1 + \alpha_2)^3} \left[T - \frac{1}{\alpha_1 + \alpha_2} (1 - \exp(-(\alpha_1 + \alpha_2)T)) \right]. \quad (19)$$

The amplitude $\mathcal{A}(\alpha_1, \alpha_2)$ is the theoretical prediction for noise correlations within single columns (in the limit where $d = |\mathbf{x} - \mathbf{y}| \rightarrow \mathbf{0}$) used in Fig. 3.

Data availability

All behavioral and electrophysiological data are archived at the Stanford Neuroscience Institute server at Stanford University. The data are available from the corresponding author upon request.

Code availability

The source code to reproduce results of this study will be made publicly available on GitHub upon publication.

References

1. McCormick, D. A., Nestvogel, D. B. & He, B. J. Neuromodulation of brain state and behavior. *Annu. Rev. Neurosci.* **43**, 391–415 (2020).
2. McGinley, M. J. *et al.* Waking state: Rapid variations modulate neural and behavioral responses. *Neuron* **87**, 1143–1161 (2015).
3. Renart, A. & Machens, C. K. Variability in neural activity and behavior. *Curr. Opin. Neurobiol.* **25**, 211–220 (2014).
4. Lee, S.-H. & Dan, Y. Neuromodulation of brain states. *Neuron* **76**, 209–222 (2012).
5. Harris, K. D. & Thiele, A. Cortical state and attention. *Nat. Rev. Neurosci.* **12**, 509–523 (2011).
6. Cohen, M. R. & Kohn, A. Measuring and interpreting neuronal correlations. *Nat. Neurosci.* **14**, 811–819 (2011).
7. Pinto, L. *et al.* Fast modulation of visual perception by basal forebrain cholinergic neurons. *Nat. Neurosci.* **16**, 1857–1863 (2013).
8. Ruff, D. A. & Cohen, M. R. Global Cognitive Factors Modulate Correlated Response Variability between V4 Neurons. *J. Neurosci.* **34**, 16408–16416 (2014).
9. Vinck, M., Batista-Brito, R., Knoblich, U. & Cardin, J. A. Arousal and Locomotion Make Distinct Contributions to Cortical Activity Patterns and Visual Encoding. *Neuron* **86**, 740–754 (2015).
10. Cohen, M. R. & Maunsell, J. H. R. Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* **12**, 1594–1600 (2009).
11. Mitchell, J. F., Sundberg, K. A. & Reynolds, J. H. Spatial Attention Decorrelates Intrinsic Activity Fluctuations in Macaque Area V4. *Neuron* **63**, 879–888 (2009).
12. Ruff, D. A. & Cohen, M. R. Attention Increases Spike Count Correlations between Visual Cortical Areas. *J. Neurosci.* **36**, 7523–7534 (2016).
13. Nandy, A. S., Nassi, J. J. & Reynolds, J. H. Laminar Organization of Attentional Modulation in Macaque Visual Area V4. *Neuron* **93**, 235–246 (2017).

14. Denfield, G. H., Ecker, A. S., Shinn, T. J., Bethge, M. & Tolias, A. S. Attentional fluctuations induce shared variability in macaque primary visual cortex. *Nat. Commun.* **9**, 1–14 (2018).
15. Abbott, L. F. & Dayan, P. The effect of correlated variability on the accuracy of a population code. *Neural Comput.* **11**, 91–101 (1999).
16. Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, 358–366 (2006).
17. Moreno-Bote, R. *et al.* Information-limiting correlations. *Nat. Neurosci.* **17**, 1410–1417 (2014).
18. Kohn, A., Coen-Cagli, R., Kanitscheider, I. & Pouget, A. Correlations and Neuronal Population Information. *Annu. Rev. Neurosci.* **39**, 237–256 (2016).
19. Nogueira, R. *et al.* The effects of population tuning and trial-by-trial variability on information encoding and behavior. *J. Neurosci.* **40**, 1066–1083 (2020).
20. Sakata, S. & Harris, K. D. Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. *Neuron* **64**, 404–418 (2009).
21. Mochol, G., Hermoso-Mendizabal, A., Sakata, S., Harris, K. D. & de la Rocha, J. Stochastic transitions into silence cause noise correlations in cortical circuits. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 3529–3534 (2015).
22. Engel, T. A. *et al.* Selective modulation of cortical state during spatial attention. *Science* **354**, 1140–1144 (2016).
23. van Kempen, J. *et al.* Top-down coordination of local cortical state during selective attention. *BioRxiv* preprint available at <https://www.biorxiv.org/content/10.1101/2020.03.26.009365v1> (2020).
24. Huang, C. *et al.* Circuit Models of Low-Dimensional Shared Variability in Cortical Networks. *Neuron* **101**, 337–348.e4 (2019).
25. Hennequin, G., Ahmadian, Y., Rubin, D. B., Lengyel, M. & Miller, K. D. The Dynamical Regime of Sensory Cortex: Stable Dynamics around a Single Stimulus-Tuned Attractor Account for Patterns of Noise Variability. *Neuron* **98**, 846–860 (2018).

26. Pettine, W. W., Steinmetz, N. A. & Moore, T. Laminar segregation of sensory coding and behavioral readout in macaque V4. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 14749–14754 (2019).
27. Jercog, D. *et al.* UP-DOWN cortical dynamics reflect state transitions in a bistable network. *ELife* **6**, e22425 (2017).
28. Glauber, R. J. Time-dependent statistics of the Ising model. *J. Math. Phys.* **4**, 294–307 (1963).
29. Rabinowitz, N. C., Goris, R. L. T. & Cohen, M. R. Attention stabilizes the shared gain of V4 populations. *Elife* **4**, e08998 (2015).
30. Smith, M. A. & Sommer, M. A. Spatial and temporal scales of neuronal correlation in visual area v4. *J. Neurosci.* **33**, 5422–5432 (2013).
31. Smith, M. A., Jia, X., Zandvakili, A. & Kohn, A. Laminar dependence of neuronal correlations in visual cortex. *J. Neurophysiol.* **109**, 940–947 (2013).
32. 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).
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. Anderson, J. C., Kennedy, H. & Martin, K. A. Pathways of attention: Synaptic relationships of frontal eye field to V4, lateral intraparietal cortex, and area 46 in macaque monkey. *J. Neurosci.* **31**, 10872–10881 (2011).
35. Rosenbaum, R., Smith, M. A., Kohn, A., Rubin, J. E. & Doiron, B. The spatial structure of correlated neuronal variability. *Nat. Neurosci.* **20**, 107–114 (2017).
36. Darshan, R., Wood, W. E., Peters, S., Leblois, A. & Hansel, D. A canonical neural mechanism for behavioral variability. *Nat. Commun.* **8**, 15415–13 (2017).
37. Sompolinsky, H., Crisanti, A. & Sommers, H. J. Chaos in random neural networks. *Phys. Rev. Lett.* **61**, 259–262 (1988).
38. Van Vreeswijk, C. & Sompolinsky, H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724–1726 (1996).
39. Renart, A. *et al.* The asynchronous state in cortical circuits. *Science* **327**, 587–590 (2010).

40. Darshan, R., Van Vreeswijk, C. & Hansel, D. Strength of Correlations in Strongly Recurrent Neuronal Networks. *Phys. Rev. X* **8**, 031072 (2018).
41. Murray, J. D. *et al.* A hierarchy of intrinsic timescales across primate cortex. *Nat. Neurosci.* **17**, 1661–1663 (2014).
42. Litwin-Kumar, A. & Doiron, B. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat. Neurosci.* **15**, 1498–1505 (2012).
43. Parga, N. Network model of spontaneous activity exhibiting synchronous transitions between up and down states. *Front. Neurosci.* **1**, 4 (2007).
44. Stringer, C. *et al.* Inhibitory control of correlated intrinsic variability in cortical networks. *ELife* **5**, e19695 (2016).
45. Holcman, D. & Tsodyks, M. The emergence of up and down states in cortical networks. *PLoS Comput. Biol.* **2**, e23 (2006).
46. Kanashiro, T., Ocker, G. K., Cohen, M. R. & Doiron, B. Attentional modulation of neuronal variability in circuit models of cortex. *Elife* **6**, 79 (2017).
47. Moore, T. & Armstrong, K. M. Selective gating of visual signals by microstimulation of frontal cortex. *Nature* **421**, 370–373 (2003).
48. Noudoost, B. & Moore, T. Control of visual cortical signals by prefrontal dopamine. *Nature* **474**, 372–375 (2011).
49. Herrero, J. L. *et al.* Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature* **454**, 1110–1114 (2008).
50. Roach, J. P., Ben-Jacob, E., Sander, L. M. & Zochowski, M. R. Formation and Dynamics of Waves in a Cortical Model of Cholinergic Modulation. *PLoS Comp. Biol.* **11**, e1004449–16 (2015).
51. Steinmetz, N. A. & Moore, T. Eye Movement Preparation Modulates Neuronal Responses in Area V4 When Dissociated from Attentional Demands. *Neuron* **83**, 496–506 (2014).
52. Lindner, B. Coherence and stochastic resonance in nonlinear dynamical systems. *PhD dissertation, Humboldt University of Berlin, Germany* (2002).

Acknowledgements

This work was supported by the Swartz Foundation (Y.S.), the NIH grant R01 EB026949 (T.A.E.), the Pershing Square Foundation (T.A.E.), the NIH grant EY014924 (T.M.).

Author Contributions

Y.S., N.A.S., K.B, T.M., and T.A.E. designed the study. N.A.S. and T.M. designed the experiments. N.A.S. performed experiments, spike sorting, and RF measurements. Y.S. and T.A.E. developed data analysis methods and mathematical models. Y.S. analyzed the data, performed analytical calculations and model simulations. Y.S., N.A.S., K.B, T.M., and T.A.E. discussed the findings and wrote the paper.

Supplementary Information

Materials and Methods

Supplementary Figures S1–S4

Supplementary Tables S1–S2

Competing interests

The authors declare no competing interests.

Figure 1. On-Off dynamics predict the magnitude of noise correlations within single cortical columns. (a) Different recording techniques sample neurons along different anatomical dimensions. A rectangular Utah multi-electrode array (top) samples from laterally separated neurons in different columns, preferentially from superficial cortical layers. A linear multi-electrode array (right) samples from neurons across all layers within single cortical columns. (b) An example trial showing spontaneous transitions between episodes of vigorous (On) and faint (Off) spiking in multi-unit activity simultaneously recorded from all layers of a single column in V4. Spikes (vertical ticks) on 16 recording channels are segmented into On (yellow) and Off (blue) episodes by the HMM. (c) Scatter plot of the variance explained by the two-phase HMM versus average noise correlation across recording sessions for two-phase (red circles) and one-phase (teal diamonds) recordings (upper panel). Stacked histogram of average noise correlations for two-phase (red) and one-phase (teal) recordings (lower panel). (d) Average noise correlations of MU in two-phase (upper panel) and one-phase (lower panel) recordings, separately for superficial and deep cortical layers, in attention (red) and control (grey) conditions. Noise correlations slightly decrease in superficial and increase in deep layers, but the overall magnitude of changes is very small.

Figure 2. A network model of interacting cortical columns. (a) The network consists of two parallel two-dimensional lattices corresponding to superficial and deep cortical layers. Each unit represents a local population of neurons within one layer of a single column. The units transition between the On (yellow) and Off (blue) phases. Bottom-up inputs from visual stimuli (grey shading) and top-down attentional inputs (red shading) target local groups of columns in the model. (b) Dynamical system modeling On-Off dynamics in single columns (left panel). The mean firing-rate variable $r(t)$ receives a recurrent self-coupling $F(r)$ and a negative feedback from the adaptation variable $a(t)$. The dynamical system is driven by a white noise $\xi(t)$, recurrent inputs from the neighboring columns $I_{\text{rec}}(t)$, and external inputs $I_{\text{stim}}(t)$ and $I_{\text{att}}(t)$. On the phase plane (right panel), the r -nullcline (grey) and a -nullcline (black) cross at the On (yellow) and Off (blue) stable fixed points. The attentional input shifts the r -nullcline (red) modulating the stability of the On and Off fixed points. (c) In single columns, the model generates stochastic On-Off transitions (left panel). The durations of On and Off episodes are irregular and exponentially distributed (right panel). The average duration (dashed lines) of On-episodes is longer in attention ($I_{\text{att}} > 0$, lower row, average $\bar{\tau}_{\text{on}} = 127$ ms, $\bar{\tau}_{\text{off}} = 83$ ms) relative to control condition ($I_{\text{att}} = 0$, upper row, average $\bar{\tau}_{\text{on}} = 102$ ms, $\bar{\tau}_{\text{off}} = 100$ ms). (d) The network generates spatiotemporal On-Off dynamics, where the On and Off phases form local spatial clusters (a single snapshot of simulated activity in the dynamical-system network is shown). The spatiotemporal pattern differs between attention (red square) and control (black square) conditions. (e) Spikes of individual neurons are modeled as inhomogeneous Poisson processes with different mean rates during the On (yellow) and Off (blue) phases generated by the network. All neurons represented by a single network unit follow the same shared On-Off sequence. (f) Model parameters are estimated by fitting the experimental data with the HMM, which provides the On-Off transition rates (α_1 and α_2 , top) and the On and Off firing rates (r_{on} and r_{off}) for each MU and SU in each recording session and task condition. Histograms show the On (yellow) and Off (blue) firing rates for MUs for an example HMM fit.

Figure 3. The model accounts for attentional modulation of correlated variability in single columns. (a) The model predicts that FF is proportional the On-Off firing-rate difference $\Delta r = r_{\text{on}} - r_{\text{off}}$. For an example recording, Δr ranges broadly across MUs (left panel), and this variation of Δr tightly corresponds with the variation of FF (right panel). (b) The model predicts that noise correlation between neurons i and j is proportional to the product $\Delta r_i \Delta r_j$. For the same example recording, the variation in $\sqrt{\Delta r_i \Delta r_j}$ corresponds with the variation of noise correlations (right panel) (c) Comparison between attention-related changes ($\Delta \text{FF} = \text{FF}_{\text{att}} - \text{FF}_{\text{ctl}}$) in FF predicted by the On-Off dynamics model (x -axis) and measured directly from the data (y -axis) for MUs. All MUs are divided in 10 equally-sized groups based on the change in their On-Off firing-rate difference between attention and control conditions ($\Delta r_{\text{att}} - \Delta r_{\text{ctl}}$, color axis). (d) Comparison between attention-related changes ($\Delta \text{NC} = \text{NC}_{\text{att}} - \text{NC}_{\text{ctl}}$) in noise correlations (NC) predicted by the On-Off dynamics model (x -axis) and measured directly from the data (y -axis) for MUs. All MU-pairs are divided in 10 equally-sized groups based on the change in the pair's On-Off firing-rate difference defined as $\sqrt{\Delta r_{\text{att},i} \Delta r_{\text{att},j}} - \sqrt{\Delta r_{\text{ctl},i} \Delta r_{\text{ctl},j}}$ (color axis).

Figure 4. Dependence of noise correlations on lateral distance. (a) Laminar recordings generally exhibit slight horizontal displacements which manifest in a systematic shift of the RFs (circles) across channels (left panel). Shift of the RFs (lines, RF contours; dots, RF centers; dva, degrees of visual angle) for an example recording (right panel). (b) In two-phase recordings, noise correlations decrease with the RF-center distance in both superficial (crimson) and deep (green) layers (dots - data points, lines - linear regression). Orange background highlights the range of short lateral distances within single or nearby columns. Purple background highlights longer lateral distances between distant columns, such as distances covered by a Utah array, which are outside the range of our laminar recordings. Error bars represent the standard error of the mean (SEM). (c) Same as b for one-phase recordings. Noise correlations do not decrease with the RF-center distance. (d) Our theory predicts that noise correlations decay with lateral distance exponentially, with the decay constant L called correlation length. Simulations of the full dynamical-system network (circles) agree with the analytical formula derived using the binary-unit network approximation (line). The model parameters α_1 , α_2 , r_{on} and r_{off} are sampled from a distribution of parameters in HMMs fitted to the data.

Figure 5. Attentional inputs modulate the efficacy of lateral interactions in the network leading to changes of the correlation length. (a) In simulations of the dynamical-system network model, noise correlations between neurons in different columns robustly decrease during attention ($I_{\text{att}} > 0$, red) relative to control ($I_{\text{att}} = 0$, grey), while noise correlations between neurons within columns change only slightly. The average reduction of noise correlations is large across columns ($\text{MI}_{\text{corr}} = -0.28$, right), but small within columns ($\text{MI}_{\text{corr}} = -0.03$, left). (b) In simulations, noise correlations decrease with excitatory ($I_{\text{att}} > 0$) and increase with inhibitory ($I_{\text{att}} < 0$) attentional inputs. In all cases, the average changes of noise correlations within columns are very small relative to sizable changes across columns. (c) Noise correlations decay faster with lateral distance in attention (red, $I_{\text{att}} > 0$) relative to control condition (grey, $I_{\text{att}} = 0$), hence the correlation length is reduced $L_{\text{att}} < L_{\text{ctl}}$. Data are shown from simulations of the full dynamical-system network (circles) and analytical calculations using the binary-network approximation (solid lines). Orange and purple backgrounds highlight the range of distances within and across columns, respectively. (d) The Off-to-On (α_1 , blue) and On-to-Off (α_2 , yellow) transition rates weakly depend on the attentional input in the dynamical-system network. (e) The effective coupling strength β steeply decreases with increasing attentional input. We estimate β from simulations of the dynamical-system network. First, we calculate the On-Off transition rates α_1 and α_2 from the average On- and Off-episode durations. Then, we estimate the correlation length L by fitting the exponential distance-dependence of noise correlations. Finally, β is computed using the analytical formula $L = \sqrt{\beta/(\alpha_1 + \alpha_2)}$.

Figure 6. Attentional changes of noise correlations depend on lateral distance. (a) The model predicts that changes of noise correlations at intermediate lateral distances are driven by changes of the correlation length L . When the correlation length decreases ($L_{\text{att}} < L_{\text{ctl}}$, left panel), the model predicts a robust reduction of noise correlations at intermediate distances (purple background), even when noise correlations at zero distance ($d = 0$, orange background) do not change. When the correlation length increases ($L_{\text{att}} > L_{\text{ctl}}$, right panel), the model predicts a robust increase of noise correlations at intermediate distances. In both cases, the spatial profile of noise-correlation changes is non-monotonic, and the greatest changes are predicted to occur at intermediate lateral distances. (b) In two-phase recordings, changes of noise correlations during attention depend on the RF-center distance. The magnitude of noise-correlation changes is vanishing at very short distances and progressively increases at longer distances. With increasing RF-center distance, noise correlations decrease in superficial (crimson) and increase in deep (green) layers (dots - data points, lines - linear regression). Error bars represent SEM. (c) Same as b for one-phase recordings. Changes of noise correlations during attention do not depend on the RF-center distance.

Figure-1

Shi et al.

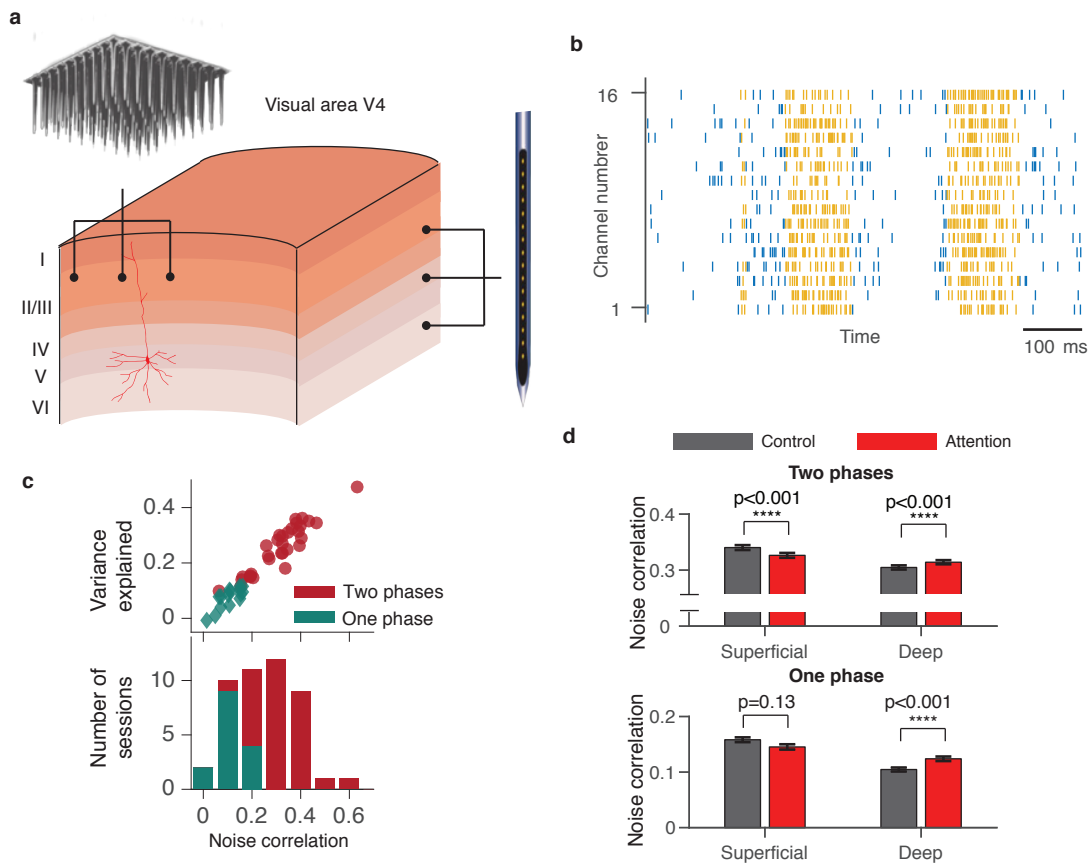


Figure-2

Shi et al.

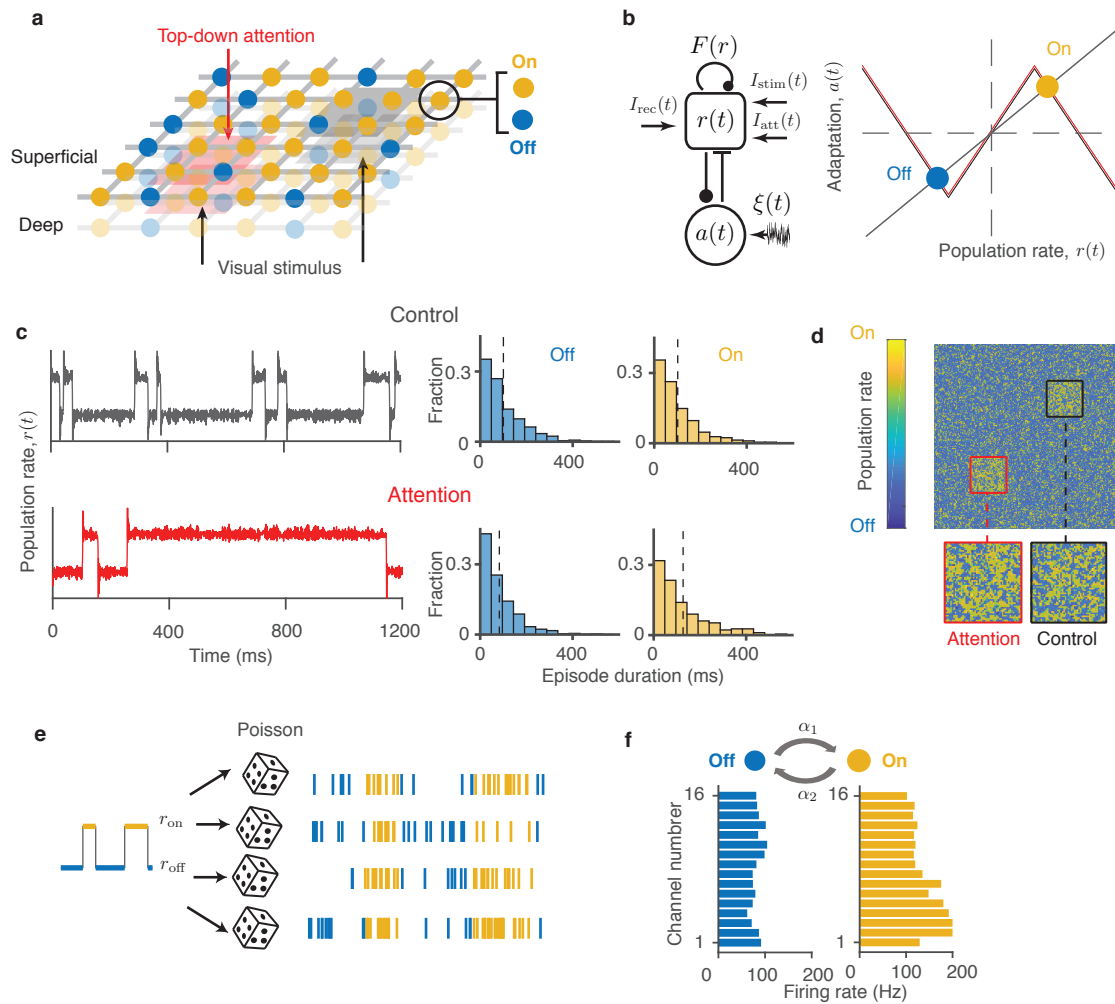


Figure-3

Shi et al.

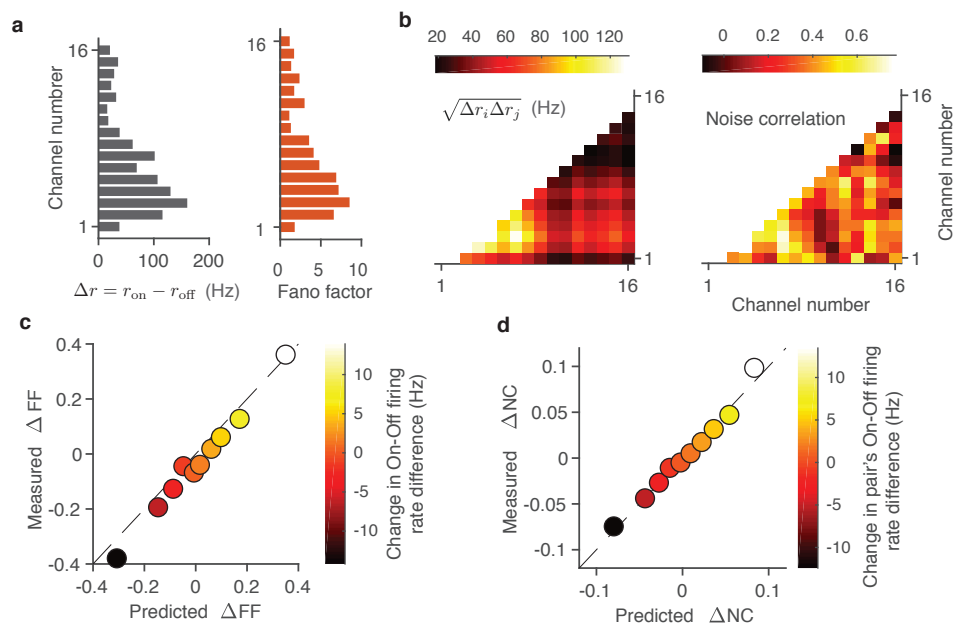


Figure-4

Shi et al.

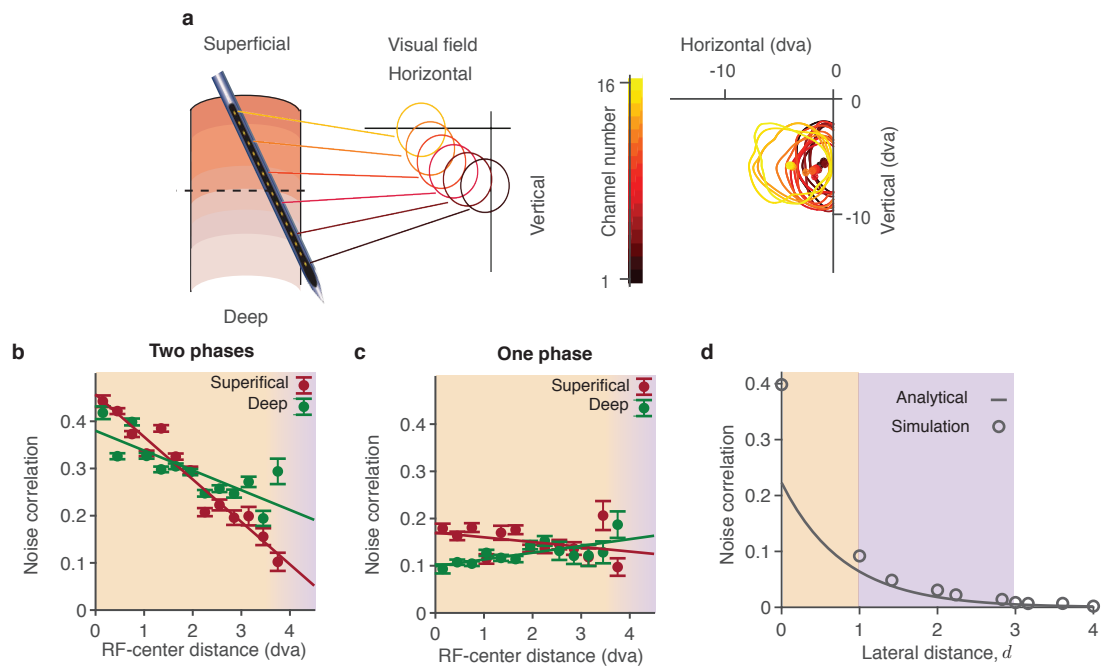


Figure-5

Shi et al.

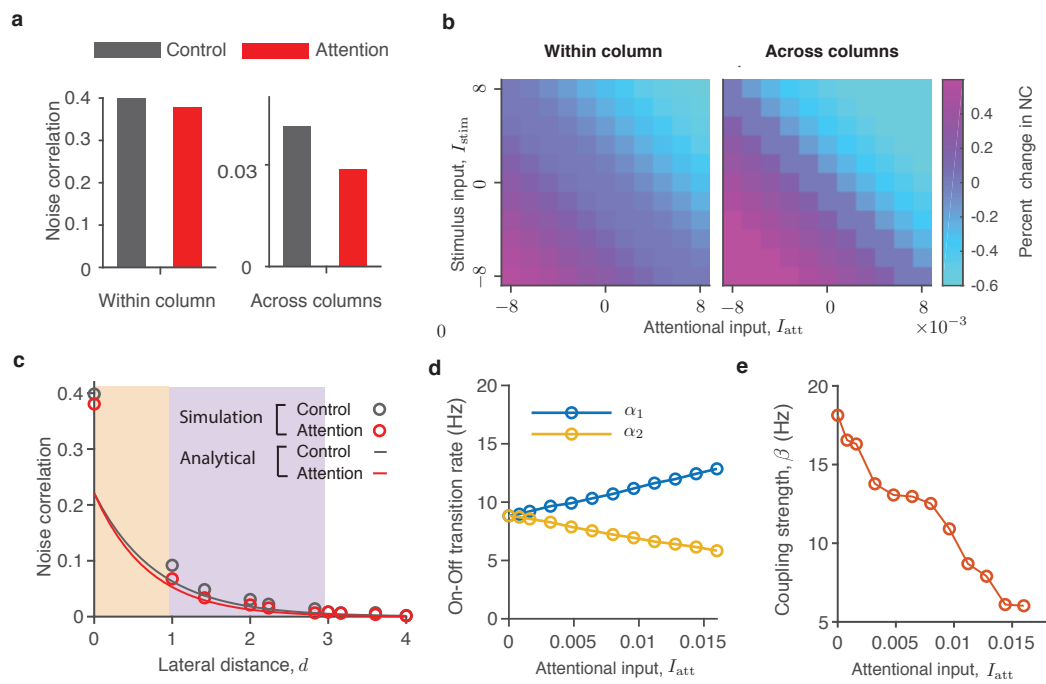


Figure-6

Shi et al.

