Non-monotonic spatial structure of interneuronal correlations in prefrontal microcircuits

Shervin Safavi [1,2,7], Abhilash Dwarakanath [1,7], Vishal Kapoor[1,2], Joachim Werner[1], Nicholas G. Hatsopoulos[3], Nikos K. Logothetis [1,4], Theofanis I. Panagiotaropoulos [1,5,6*]

[1]Department of Physiology of Cognitive Processes, Max Planck Institute for Biological Cybernetics, Tübingen 72076, Germany

[2] IMPRS for Cognitive and Systems Neuroscience, University of Tübingen, Tübingen 72074, Germany

[3]Department of Organismal Biology and Anatomy, University of Chicago, USA

[4]Division of Imaging Science and Biomedical Engineering, University of Manchester, Manchester M13 9PT, UK

[5] Centre for Systems Neuroscience, University of Leicester, Leicester LE1 7QR, UK

[6] Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, UK

[7] These authors contributed equally to this work

*Correspondence: theofanis.panagiotaropoulos@tuebingen.mpg.de

Summary

Correlated fluctuations of single neuron discharges, on a mesoscopic scale, decrease as a function of lateral distance in early sensory cortices, reflecting a rapid spatial decay of lateral connection probability and excitation. However, spatial periodicities in horizontal connectivity and associational input as well as an enhanced probability of lateral excitatory connections in the association cortex could theoretically result in non-monotonic correlation structures. Here we show such a spatially non-monotonic correlation structure, characterized by significantly positive long-range correlations, in the inferior convexity of the macaque prefrontal cortex. This functional connectivity kernel was more pronounced during wakefulness than anesthesia and could be largely attributed to the spatial pattern of correlated variability between functionally similar neurons during structured visual stimulation. These results suggest that the spatial decay of lateral functional connectivity is not a common organizational principle of neocortical microcircuits. A non-monotonic correlation structure could reflect a critical topological feature of prefrontal microcircuits, facilitating their role in integrative processes.

eTOC Blurb

Safavi et al. report a non-monotonic and long-range spatial correlation structure in the lateral prefrontal cortex.

Highlights

- Non-monotonic spatial structure of interneuronal correlations in prefrontal cortex
- Non-monotonic spatial structure stronger during wakefulness than anesthesia
- Strong and positive long-range correlations between prefrontal neurons
- Non-monotonicity stronger for functionally similar neurons

Introduction

The intra-areal connectivity patterns of neural populations in the mammalian neocortex frequently repeat across cortical areas (Douglas and Martin, 2004, 2007; Douglas et al., 1989; Harris and Mrsic-Flogel, 2013). Such canonical rules with general validity are important in understanding basic organizational principles and ensemble computations in cortical networks (Douglas and Martin, 2004; Douglas et al., 1989; Harris and Mrsic-Flogel, 2013; Miller, 2016). Nevertheless, identifying deviations from these rules between sensory and higher order, association cortical areas could reveal properties leading to cortical network specialization and higher cognitive functions (Douglas and Martin, 2004; Douglas et al., 1989; Harris and Mrsic-Flogel, 2013; Miller, 2016; Murray et al., 2014).

The spatial structure of intra-areal functional connectivity is frequently inferred by measuring the trial-by-trial correlated variability of neuronal discharges (spike count correlations) (Cohen and Kohn, 2011). One of the most well-established properties (a canonical feature) of intra-areal, mesoscopic, functional connectivity is a so called limited-range correlation structure, reflecting a monotonic decrease of spike count correlations as a function of spatial distance and tuning similarity (Bair et al., 2001; Chelaru and Dragoi, 2016; Cohen and Kohn, 2011; Constantinidis and Goldman-Rakic, 2002; Denman and Contreras, 2014; Ecker et al., 2014; Rothschild et al., 2010; Schulz et al., 2015; Smith and Kohn, 2008; Smith and Sommer, 2013; Sompolinsky et al., 2001). However, this distance-dependent decrease of correlations has been almost exclusively derived from recordings in primary sensory cortical areas or inferred from recordings in the PFC with various constraints like a rather limited scale (Constantinidis et al., 2001) (see also Discussion section). As a result, it is currently unclear whether known differences in the structure of anatomical connectivity across the cortical

hierarchy could also give rise to different spatial patterns of functional connectivity (Amir et al., 1993; Elston, 2003; Gochin et al., 1991; Lund et al., 1993).

Specifically, the rapid spatial decay of correlations in sensory cortex is widely assumed to reflect a similar rapid decay in lateral anatomical connectivity and excitation (Martin et al., 2014). In early visual cortical areas, correlations rapidly decrease as a function of distance (Bair et al., 2001; Ecker et al., 2014; Smith and Kohn, 2008; Smith and Sommer, 2013; but also see (Ecker et al., 2010; Rosenbaum et al., 2017)) in a manner that closely reflects anatomical findings about the limited spread and density of intrinsic lateral connections (Amir et al., 1993; Angelucci et al., 2002; Kritzer and Goldman-Rakic, 1995; Tanigawa et al., 2005; Voges et al., 2010). However, lateral connections are significantly expanded in later stages of the cortical hierarchy, like the prefrontal cortex (PFC) (Amir et al., 1993; Fujita and Fujita, 1996; Kritzer and Goldman-Rakic, 1995; Levitt et al., 1993; Lund et al., 1993; Tanigawa et al., 2005). In this higher-order association area, lateral connections commonly extend to distances up to 7-8mm (Kritzer and Goldman-Rakic, 1995; Levitt et al., 1993; Tanigawa et al., 2005) while patches of connected populations are both larger and more distant from each other compared to sensory cortex (Pucak et al., 1996; Tanigawa et al., 2005). Although horizontal axons in macaque V1 can extend up to 4mm, they do not form clear patches, and for distances of 2-3mm laterally to the injection patch border, only a small number of cells are labeled in comparison to higher-order areas (Amir et al., 1993; Angelucci et al., 2002; Schwartz and Goldman-Rakic, 1984; Tanigawa et al., 2005; Yoshioka et al., 1996). In addition to the more extended intrinsic lateral connectivity, associational input from other cortical areas to the PFC also forms stripes with an average distance of 1.5mm and contributes to the spatial periodicities in lateral organization (Goldman-Rakic and Schwartz, 1982). Finally, the proportion of lateral excitatory connections is

higher in the PFC (95%) compared to V1 (75%) (Melchitzky et al., 2001).

Whether these significant differences in the structural architecture of the PFC compared to early sensory areas also result in a distinct spatial pattern of functional connectivity is currently unknown. Intuitively, higher probability of long-range lateral excitatory connections and stripe-like connectivity patterns could give rise to strong spike count correlations across local and spatially remote populations with weaker correlations for populations in intermediate distances. To address this question, we recorded simultaneously the activity of large neural populations in the inferior convexity of the macaque PFC during both anesthetized and awake states using multi-electrode Utah arrays (Maynard et al., 1997). In both anesthetized and awake states the spatial pattern of pairwise correlated variability was non-monotonic with significantly positive long-range correlations. A major source of non-monotonicity could be attributed to the spatial pattern of correlated variability between functionally similar neurons.

Results

We used multi-electrode Utah arrays (4x4mm, 10x10 electrodes, inter-electrode distance 400μm, electrode length 1mm, Figure 1A) to record spiking activity from the inferior convexity of the ventrolateral prefrontal cortex (vlPFC) during repeated visual stimulation with movie clips in two anesthetized macaque monkeys (Figure 1B) and with sinusoidal gratings, drifting in 8 different directions, in two awake behaving macaques (Figure 1C). To evaluate the effect of structured visual input on correlated variability, we contrasted periods of visual stimulation to intertrial as well as spontaneous activity (long periods of neural activity without any task demands). Both anesthetized and awake state recordings resulted in the simultaneous monitoring

of multiple, well isolated single units that remained stable for hours of recording (Figure 1D). On average, in each dataset, we recorded from 103 ± 16 (mean ± S.E.M) single units and 5305 ± 1681 pairs during anesthesia (Figure S1A), and 81 ± 12 single units and 3281 ± 980 pairs during the awake state recordings (Figure S1B).

Spatial structure of correlated variability during anesthesia and wakefulness

It has been repeatedly shown that correlated variability of spike counts in early sensory, especially visual, areas in different species decreases as a function of lateral distance with strong interactions for proximal and progressively weaker interactions for distal (up to 4mm) neurons (Schulz et al., 2015; Smith and Kohn, 2008; Smith and Sommer, 2013). We investigated the same relationship between spike count correlations (r_{sc}) and lateral distance up to 4mm in the vlPFC.

Visual stimulation with movie clips during anesthesia or with drifting gratings during wakefulness, gave rise to a spatial pattern in the structure of correlated variability that was fundamentally different compared to early sensory areas: Strong and positive long-range (> 2.5mm) correlations that were comparable to the average magnitude of local (up to 0.5mm) correlations and significantly weaker correlations for intermediate distances (red curves in Figure 2A and 2B and Figure S1A and S1B).

We evaluated the statistical significance of non-monotonicity and long-range correlations by comparing the distributions of pairwise correlations in populations recorded from nearby (0.5mm), intermediate (2.5mm) and distant (3.5-4mm) sites during visual stimulation. Average correlated variability between neurons located in intermediate distances was significantly lower compared to nearby neurons in both anesthetized ($\overline{r_{sc}}^{0.5mm}$ = 0.0675 ± 0.0016 vs. $\overline{r_{sc}}^{2.5mm}$ =

0.0520 ± 0.0010; $p < 10^{-10}$, Wilcoxon rank-sum test, Figure 2A red curve and Figure S2A) and awake recordings ($\overline{r_{sc}}^{0.5mm}$ = 0.0193 ± 0.0026 vs. $\overline{r_{sc}}^{2.5mm}$ = 0.0046 ± 0.0015; p = 0.0038, Figure 2B red curve and Figure S2D). Following this minimum, correlations during anesthesia significantly increased from 2.5mm to 3mm ($\overline{r_{sc}}^{2.5mm}$ = 0.0520 ± 0.0010 vs. $\overline{r_{sc}}^{3mm}$ = 0.0592 ± 0.0012; p = 0.012) and 3.5mm ($\overline{r_{sc}}^{3.5mm}$ = 0.0605 ± 0.0018; p = 0.0040) but not for 4mm (p = 0.4123). A similar increase in correlated variability for progressively more distant populations was also observed in the awake state, where correlations significantly increased from 2.5 to both 3.5mm ($\overline{r_{sc}}^{2.5mm}$ = 0.0046 ± 0.0015 vs. $\overline{r_{sc}}^{3.5mm}$ = 0.0131 ± 0.0024; p = 0.0093) and 4mm ($\overline{r_{sc}}^{2.5mm}$ = 0.0046 ± 0.0015 vs. $\overline{r_{sc}}^{4mm}$ = 0.0190 ± 0.0041; p = 0.0110).

In the awake state recordings, the average magnitude of correlations for distant populations, located 3.5-4mm apart, was not different from the respective magnitude for nearby pairs ($\overline{r_{sc}}^{0.5mm}$ = 0.0193 ± 0.0026 vs. $\overline{r_{sc}}^{3.5mm}$ = 0.0131 ± 0.0024; p > 0.6 and $\overline{r_{sc}}^{4mm}$ = 0.0190 ± 0.0041; p > 0.7, Figure 2B red curve and Figure S2D). In addition, both local and distant average correlations were significantly positive (p < 0.005, t-test). However, in the anesthetized recordings, despite the significant increase of correlations for distant neurons compared to intermediate distances, long-range correlations were significantly lower compared to nearby neurons ($\overline{r_{sc}}^{0.5mm}$ = 0.075 ± 0.0016 vs. $\overline{r_{sc}}^{3.5mm}$ = 0.0605 ± 0.0018; p = 0.0056 and $\overline{r_{sc}}^{4mm}$ = 0.0544 ± 0.0027; p = 0.0008, Figure 2A red curve and Figure S2A), suggesting that anesthesia induced fluctuations had a non-homogeneous impact on the spatial structure of correlations. Such non-homogeneous, state-dependent weighting on the spatial structure of correlations has been reported in previous studies of primary visual cortex as well (Ecker et al., 2014).

The decrease in correlations from 0.5 to 2.5mm and the increase from 2.5 to 3.5 or 4mm was observed in both anesthetized and awake states. However, in the awake state recordings, we

also observed an additional pronounced peak at 2mm (Figure 2B and Figure S2D). Lack of this peak at intermediate distances in our anesthetized recordings is compatible with other studies performed during anesthesia and provide further evidence for a non-homogeneous, state-dependent weighting on the spatial structure of correlated variability (Ecker et al., 2014; Rosenbaum et al., 2017; Smith and Kohn, 2008; Smith and Sommer, 2013). Indeed, these common features in the spatial structure of correlated activity across anesthetized ($\overline{r_{sc}}^{anesth}$) and awake states ($\overline{r_{sc}}^{awake}$) were observed despite significant differences in the average magnitude of correlations ($\overline{r_{sc}}^{anesth}$ = 0.0574 ± 3 x 10^{-4} (mean ± S.E.M.) vs. $\overline{r_{sc}}^{awake}$ = 0.0122 ± 4 x 10^{-4}; p = 0, Wilcoxon rank-sum test, Figure 3A). Despite being very close to zero, average correlations during visual stimulation were significantly positive during the awake state ($p < 10^{-33}$; t-test).

Visual stimulation shapes the spatial structure of correlated variability

We evaluated the impact of structured visual stimulation on the spatial pattern of correlated variability by comparing correlations during visual stimulation with movie clips (during anesthesia) or drifting sinusoidal gratings (during wakefulness) to the respective pattern during intertrial and spontaneous activity periods. Compared to periods of intertrial activity visual stimulation resulted in a rather weak increase of correlated variability in both anesthetized ($\overline{r_{sc}}^{visual}$ = 0.0574 ± 3 x 10^{-4} vs. $\overline{r_{sc}}^{intertrial}$ = 0.0554 ± 3 x 10^{-4}; $p < 10^{-3}$; Figure 3B) and awake recordings ($\overline{r_{sc}}^{visual}$ = 0.0122 ± 4 x 10^{-4} vs. $\overline{r_{sc}}^{intertrial}$ = 0.0090 ± 4 x 10^{-4} ; p = 0.1096, Figure 3C).

Despite a negligible impact on the average magnitude of correlations, visual stimulation in the awake state significantly shaped a spatially inhomogeneous, non-monotonic structure of correlated variability. In striking contrast to the significant differences observed for the same

lateral distances during visual stimulation, we found that correlations during the intertrial period were not different between local and intermediate ($\overline{r_{sc}}^{0.5mm}$ = 0.0037 ± 0.0024 vs. $\overline{r_{sc}}^{2.5mm}$ = 0.0072 ± 0.0016; p > 0.3, Figure 2B black curve and Figure S2B) or intermediate and distant populations ($\overline{r_{sc}}^{2.5mm}$ = 0.0072 ± 0.0016 vs. $\overline{r_{sc}}^{3.5mm}$ = 0.0107 ± 0.0020; p > 0.2 and vs. $\overline{r_{sc}}^{4mm}$ = 0.0128 ± 0.0042 , p > 0.2, Figure 2B black curve and Figure S2B). Spatially homogeneous correlations were also observed during periods without any structured visual input or task engagement, in data collected during spontaneous activity (Figure 2B blue curve). In these epochs, we also found no difference between local and intermediate correlations ($\overline{r_{sc}}^{0.5mm}$ = 0.0110 ± 0.0015 vs. $\overline{r_{sc}}^{2.5mm}$ = 0.0094 ± 0.0010; p > 0.7) and very similar correlations between intermediate and distant populations ($\overline{r_{sc}}^{2.5mm}$ = 0.0094 ± 0.0010 vs. $\overline{r_{sc}}^{3.5mm}$ = 0.0128 ± 0.0014; p = 0.047 and vs. $\overline{r_{sc}}^{4mm}$ = 0.0099 ± 0.0019; p > 0.5, Figure 2B blue curve).

We quantified the magnitude of spatial inhomogeneity in the structure of correlations across different conditions and states (see Experimental Procedures). A clear difference in the rate of changes in correlated variability was observed in awake state recordings (Figure 4A), where visual stimulation resulted in the strongest spatial variability and intertrial activity in the weakest (almost constant average correlation as a function of lateral distance). A similar spatial variability was also observed under anesthesia (Figure 4B), however, the average rate of change was comparable across the two conditions of visual stimulation and intertrial, but different during spontaneous activity. The difference in the structure of functional connectivity between visual stimulation and intertrial periods across anesthesia and awake states could be attributed to the lack of saccadic eye movements in intertrial periods during anesthesia. Saccadic eye movements reset visual perception (Paradiso et al., 2012) and their absence could create a persistent network state, showing no reset, resulting in very similar pattern of correlations during

visual stimulation and intertrial periods.

These results suggest that the spatial structure of correlated variability in PFC is inhomogeneous. The magnitude of inhomogeneity depended not only on the variation of global states such as wakefulness or anesthesia but most importantly on behavioural demands i.e. visual stimulation, intertrial (anticipation of the succeeding trial) or spontaneous activity (no behavioural load). Although traces of inhomogeneity in the spatial structure of correlations were observed during spontaneous activity or intertrial periods, structured visual stimulation during the awake state appeared to result in the strongest spatial inhomogeneity in the correlation structure.

Prevalence of non-monotonic spatial structure in functionally similar populations

Horizontal connectivity in PFC has been hypothesized to preferentially target neurons with functional similarities (e.g. similar spatial tuning), similar to iso-orientation columns in the visual cortex (Goldman-Rakic, 1995). Therefore, we next examined whether the source of this non-monotonicity could be traced to populations of neurons that were modulated similarly by visual input. First, tuning functions for each recorded unit were obtained based upon the discharge response to sinusoidal gratings drifting in eight different directions (Figure 5). The correlation between tuning functions (signal correlation, r_{signal}), provided a measure of functional similarity among the recorded pairs (see Experimental Procedures for more details).

We analysed the relationship between the spatial structure of functional connectivity and functional similarity of pairwise responses (i.e. signal correlations). The non-monotonic correlation structure was stronger for pairs of functionally similar neurons ($r_{signal} > 0.1$) during visual stimulation (Figure 6A). For $r_{signal} > 0.1$, correlations in intermediate distances were

significantly lower compared to nearby neurons ($\overline{r_{sc}}^{0.5mm}$ = 0.0300 ± 0.0046 vs. $\overline{r_{sc}}^{2.5mm}$ = 0.0117 ± 0.0027; p = 0.0390) while correlations significantly increased for longer distances ($\overline{r_{sc}}^{2.5mm}$ = 0.0117 ± 0.0027 vs. $\overline{r_{sc}}^{4mm}$ = 0.0298 ± 0.0069; p = 0.0418) (Figure 6B). Furthermore, the strength of both local and long-range correlated variability for positive signal correlations was significantly positive (t-test; p < 0.005) and not different between local and distant populations ($\overline{r_{sc}}^{0.5mm}$ = 0.0300 ± 0.0046 vs. $\overline{r_{sc}}^{3.5mm}$ = 0.0232 ± 0.0042; p > 0.7 and vs. $\overline{r_{sc}}^{4mm}$ = 0.0298 ± 0.0069; p > 0.6). For populations with negative signal correlations (r_{signal} < -0.1), the strength of correlated variability between nearby and intermediate neurons was not significantly different ($\overline{r_{sc}}^{0.5mm}$ = 0.0039 ± 0.0040 vs. $\overline{r_{sc}}^{2.5mm}$ = - 0.0029 ± 0.0025; p = 0.3530, Figure 6B) while correlations between neurons in intermediate and distant locations were also very similar ($\overline{r_{sc}}^{2.5mm}$ = - 0.0029 ± 0.0025 vs. $\overline{r_{sc}}^{3.5mm}$ = 0.0040 ± 0.0040; p = 0.1001; and vs. $\overline{r_{sc}}^{4mm}$ = 0.0051 ± 0.0065; p = 0.2627).

During intertrial periods correlated variability of functionally similar neurons (functional similarity estimated during the visual stimulation period) was more homogeneous although statistically non-significant traces of non-monotonicity could still be detected (Figure 6C,D). For positive signal correlations (r_{signal} > 0.1) during visual stimulation , the strength of correlated variability between nearby and intermediate neurons was not significantly different ($\overline{r_{sc}}^{0.5mm}$ = 0.0059 ± 0.0043 vs. $\overline{r_{sc}}^{2.5mm}$ = 0.0118 ± 0.0027; p = 0.4448, Figure 6D) and correlations between neurons in intermediate and distant locations were also very similar ($\overline{r_{sc}}^{2.5mm}$ = 0.0118 ± 0.0027 vs. $\overline{r_{sc}}^{3.5mm}$ = 0.0055 ± 0.0040; p = 0.3556; and vs. $\overline{r_{sc}}^{4mm}$ = 0.0106 ± 0.0071; p = 0.7176). Progressively higher and lower thresholds for signal correlation, resulting in sampling populations with stronger and weaker functional similarity respectively, did not qualitatively change these effects (Figure S3).

Since the main feature of a spatially non-monotonic correlation structure in the population of functionally similar neurons was a strengthening of short-range and long-range interactions (Figure 6B), we also compared correlated variability across lateral distance between the subpopulations exhibiting positive (green trace) and negative (blue trace) signal correlations (Figure 6B). The magnitude of the correlation coefficient was comparable only for intermediate lateral distances (i.e. 2 and 3mm, $p > 0.08$). In contrast, for local and long-range interactions correlations were significantly higher for functionally similar neurons ($p < 0.05$). Examples of pairwise neuronal responses from neurons with similar signal correlations and sampled from short, intermediate and long lateral distances are presented in Figure 6E-G.

These results suggest that spatial inhomogeneities in the functional architecture of the PFC arise from strong local and long-range lateral interactions between functionally similar neurons, which are particularly pronounced during structured visual stimulation in the awake state.

Discussion

Spatial structure of prefrontal correlated variability and relationship to anatomical structure

Spatial decay in the strength of spike count correlations on a mesoscopic scale, up to 4 mm of lateral distance, is largely considered a canonical feature of functional connectivity. Our results suggest that this spatial decay is not observed in the PFC since nearby and distant neurons are correlated to the same degree, thus reflecting a fundamentally different lateral functional connectivity structure compared to primary sensory areas like V1 (Ecker et al., 2010; Rosenbaum et al., 2017; Smith and Kohn, 2008; Smith and Sommer, 2013). Such a functional connectivity pattern is likely to directly reflect the underlying anatomical organization of prefrontal neural populations into spatially distributed clusters connected through local and long-range excitatory collaterals (Kritzer and Goldman-Rakic, 1995; Melchitzky et al., 2001; Pucak et al., 1996). Indeed, in awake state recordings, the spatially inhomogeneous correlation pattern reflects bumps of ~ 1-2 mm width (Figure 2B) which is closely matching the spatial distribution of laterally labeled stripes of neuronal assemblies (~0.8-1.5mm) in supragranular prefrontal layers (Kritzer & Goldman-Rakic 1995; Pucak et al. 1996).

Although purely anatomical methods cannot identify functional similarities across connected populations (and vice versa), an influential hypothesis of structural connectivity in the PFC assumes that long-range excitatory collaterals target clusters of neurons with similar functional preference, like spatial tuning (Goldman-Rakic, 1995). Our results provide experimental evidence supporting this hypothesis since correlated variability of functionally similar neurons was the major source of spatial inhomogeneities, on a spatial scale that closely matches the anatomical estimates of periodicities in lateral connections and associational input.

Interestingly, despite a columnar structure of orientation preference in V1 (Hubel et al. 1978), correlated variability is significantly lower for distant populations, potentially reflecting a much weaker influence of lateral connections. Although a definite answer to the exact relationship between structural and functional connectivity in the PFC could be provided in the future from functional anatomy techniques, the spatial scale consistency across anatomical and functional connectivity measures seems to suggest that indeed, structural connectivity is likely to cluster functionally similar prefrontal populations into local and distant functionally connected ensembles.

The implications of functionally interacting, spatially distributed clusters of functionally similar populations are currently not clear. Nevertheless, understanding the topology of lateral functional connectivity in the PFC could provide further constraints on recurrent neural network (RNN) models that were recently employed to simulate context dependent decision making from prefrontal neural populations (Mante et al., 2013). Implementing biologically inspired constraints like anatomical and functional inhomogeneities through weighted node interactions in the spatial structure of these models that are currently either fully or sparsely and randomly connected, could further refine RNNs (Sussillo, 2014). Moreover, such realistic, empirically estimated connectivity kernels could also be employed in neural field modelling that model a variety of dynamic processes in neuronal networks.

Spatial structure of prefrontal correlated variability and integrative processing

PFC is a central subnetwork having a crucial role in cognitive computations due to an increase in the integrative aspect of information processing in higher order cortical areas (Miller and Cohen, 2001; Modha and Singh, 2010). This progressive increase in integrative functions across the

cortical hierarchy was recently suggested to be mediated by a similar hierarchy in the timescales of intrinsic fluctuations that arise due to systematic changes in the anatomical structure, like heterogeneous connectivity of local circuitry (Chaudhuri et al., 2015; Chen et al., 2015; Murray et al., 2014). A non-monotonic spatial structure of correlated variability differentiates prefrontal functional connectivity from primary sensory areas and could therefore be relevant to the emergence of prefrontal-specific timescales (Chaudhuri et al., 2014, 2015; Chen et al., 2015; Murray et al., 2014). Network topology was recently suggested to affect timescales since physical distance between connected nodes was shown to increase as time scale lengthened (Timme et al., 2014).

From a graph-theoretical perspective, a spatially inhomogeneous connectivity profile, combining strong local and long-range functional connectivity, similar to what we observed in the PFC for functionally similar populations, could reflect a network with shorter average path length and higher average clustering coefficient compared to a network with monotonically decreasing correlations and/or uncorrelated long-range functional connectivity (like V1) (Bullmore and Sporns, 2009). These topological features are known to facilitate efficient integrative processing (Latora and Marchiori, 2001, 2003) and could reflect a fundamental characteristic of laterally organized prefrontal microcircuits compared to primary visual cortex where despite positive local correlations long-range activity on the same spatial scale is uncorrelated (Rosenbaum et al. 2017).

Comparison with previous studies of correlated variability in the prefrontal cortex

Experimental constraints prevented previous studies in dorsolateral areas of the PFC, around the principal sulcus, from capturing a non-monotonic correlation structure (Constantinidis and

Goldman-Rakic, 2002; Sakurai and Takahashi, 2006; Tsujimoto et al., 2008). These studies were constrained by a maximum inter-electrode distance of 1 mm, and our findings up to this distance are indeed consistent, showing a decrease in correlations up to 1 mm.

A number of other factors might also have prevented previous studies that used Utah arrays in other areas of the PFC to capture the non-monotonic spatial structure of correlations we report here. Firstly, it is likely that the non-monotonic structure is specific for the this particular region of PFC, i.e. vlPFC, since none of these studies involved recordings in the vlPFC but rather in area 8a (i.e. the frontal eye fields) (Kiani et al., 2015; Leavitt et al., 2013). The source of this region-specific discrepancy between our results and previous studies (Kiani et al., 2015; Leavitt et al., 2013) could be potentially traced to differences in the involvement of various prefrontal regions in visual processing. For example, the probability of finding feature selective neurons (e.g. direction selective neurons) may be higher in the vlPFC compared to area 8a (Hussar and Pasternak, 2009). Since our data validated the spatially non-monotonic correlation structure during visual stimulation with movie clips and direction of motion, the lack of a similar spatial structure in the frontal eye fields could be due to its differential functional role .

Leavitt et al. (2013) recorded using 4x4mm Utah arrays in area 8A and found a monotonically decreasing correlation structure. However, hardware limitations allowed them to record simultaneously from blocks of only 32 electrodes each time, limiting the spatial coverage that would prevent an extensive examination of the potential spatial anisotropy in area 8A. Kiani et al. (2015), using the same electrode arrays, recorded simultaneously from all 96 electrodes and also reported a monotonic decrease of correlations for multi-unit activity (MUA) for distances up to 4mm. However, the length of electrodes was 1.5mm compared to the 1mm length used in our recordings. Therefore, the monotonically decreasing correlations might be due to layer-specific

effects as previously reported in primary visual cortex (Hansen et al., 2012; Smith et al., 2013).

Comparison with previous studies of correlated variability in primary visual cortex

Rosenbaum et al. (2017) recently provided evidence for a non-monotonic correlation structure in superficial layers (L2/3) of primary visual cortex without strong long-range correlations. In particular, they reanalyzed data collected with Utah arrays during anesthesia and, only after removing the effect of latent shared variability, found that nearby neurons were weakly but significantly correlated, neurons at intermediate distances were negatively correlated and distant neurons were uncorrelated (r_{sc} not different from 0).

There are some major differences between this study and our results from prefrontal recordings on the same spatial scale. Firstly, the average correlation coefficient for distant (3-4 mm apart) neurons in these V1 recordings was not different from zero, which implies an absence of correlation rather than weak correlation between distant populations. In striking contrast, the average magnitude of long-range correlations for the same distance in the awake PFC recordings was a) significantly positive and b) comparable to the magnitude of correlations for nearby neurons. This suggests that long-range (3-4 mm) functional connectivity in PFC is stronger and in fact results in significant long-range correlations compared to the primary visual cortex where, despite a weak non-monotonicity, the average correlated variability between distant neurons is not different from zero.

The second and more important difference pertains to the conditions under which the non-monotonic structure was detectable. The Rosenbaum et al. results in V1 suggested an underlying non-monotonic functional connectivity that was washed out by the strong modulatory effects of global state fluctuations observed during anesthesia (Ecker et al. 2014). Specifically,

the non-monotonic correlation structure was revealed only after removing the effect of global latent fluctuations via Gaussian Process Factor Analysis (GPFA). This suggests that a non-monotonic structure in the primary visual cortex should be directly detectable in data recorded from awake animals where the anesthesia- induced global fluctuations are absent. However, to the best of our knowledge, until now there is no direct experimental evidence in awake V1 recordings. In contrast, Ecker et al. (2010) found a flat correlation structure in awake V1 recordings using tetrode arrays, which was also revealed after removing latent fluctuations from anesthetized recordings using the same technique (Ecker et al., 2014). Regardless of the underlying reason for this discrepancy (e.g. layer specificity of the effect or number of samples), our recordings in the PFC provide the first direct evidence for a non-monotonic, long-range correlation structure during wakefulness, without the need for removing latent sources of covariance, i.e. without application of GPFA or any other similar tool involving theoretical assumptions like stationarity of responses or the number of latent factors that contributes in driving correlated variability.

Conclusion

Overall, our results suggest that the mesoscopic functional connectivity architecture of vlPFC is fundamentally different compared to early sensory cortices, like V1. Correlated variability in the vlPFC is spatially non-monotonic and a major source of non-monotonicity is the spatial pattern of correlations between neurons with similar functional properties. A non-monotonic functional connectivity profile with strong long-range interactions might reflect the underlying machinery for large-scale coordination of distributed information processing in the prefrontal cortex.

Experimental Procedures

Electrophysiological recordings

Extracellular electrophysiological recordings were performed in the inferior convexity of the lateral PFC of 2 anesthetized and 2 awake adult, male rhesus macaques (*Macaca mulatta*) using Utah microelectrode arrays (Blackrock Microsystems, Salt Lake City, Utah USA; (Maynard et al., 1997). The array (4x4mm with a 10 by 10 electrode configuration and inter-electrode distance of 400μm) was placed 1 - 2 millimeters anterior to the bank of the arcuate sulcus and below the ventral bank of the principal sulcus, thus covering a large part of the inferior convexity in the ventrolateral PFC (Figure 1A). For the awake experiments, monkeys were implanted with form-specific titanium head posts on the cranium after modelling the skull based on an anatomical MRI scan acquired in a vertical 7T scanner with a 60cm diameter bore (Biospec 47/40c; Bruker Medical, Ettlingen, Germany). The methods for surgical preparation and anesthesia have been described in previous studies (Belitski et al., 2008; Logothetis et al., 2002, 1999). All experiments were approved by the local authorities (Regierungspräsidium) and were in full compliance with the guidelines of the European Community (EUVD 86/609/EEC) for the care and use of laboratory animals.

Data acquisition and spike sorting

Broadband neural signals (0.1-32 kHz in the anesthetized recordings, and 0.1-30 kHz in the awake recordings) were recorded using a Neuralynx (Digital Lynx) data acquisition system for the anesthetized recordings and Neural Signal Processors (Blackrock Microsystems) for the awake recordings.

In the anesthetized data, to detect spiking activity we first band–pass filtered (0.6-5.8 kHz) the broadband raw signal using a minimum–order finite impulse response (FIR) filter (Rabiner et al., 1975) with 65dB attenuation in the stop–bands and less than 0.002dB ripple within the pass–band. A gaussian distribution was fit to randomly selected chunks of the filtered signal to compute the noise variance, and the amplitude threshold for spike detection was set to 5 times the computed variance. Spike events with inter–spike intervals less than a refractory period of 0.5ms were eliminated. Those events that satisfied the threshold and refractory period criteria were kept for spike sorting.

In the awake experiments, broadband data were filtered between 0.3-3 kHz using a 2^{nd} order Butterworth filter. The amplitude for spike detection was set to 5 times the median absolute deviation (Quiroga et al., 2004). The criterion for rejection of spikes was the same as described above. All the collected spikes were aligned to the minimum. For spike sorting, 1.5 ms around the peak, i.e. 45 samples were extracted.

Automatic clustering to detect putative single neurons in both the awake and anesthetized data was achieved by a Split and Merge Expectation-Maximisation (SMEM) algorithm that fits a mixture of Gaussians to the spike feature data which consisted of the first three principal components. For the anesthetized data, the SMEM algorithm by Ueda et al was used (Ueda et al., 2000). Details of the spike sorting method used in this study have been described in other papers using tetrodes (Ecker et al., 2010; Tolias et al., 2007). For the awake data, the KlustaKwik algorithm (Harris et al., 2000; Kadir et al., 2014) was employed. The spike sorting procedure was finalized in both cases through visual inspection using the program Klusters (Hazan et al., 2006).

Visual stimulation

In anesthetized recordings, full-field visual stimulation of 640 × 480 resolution with 24-bit true colour at 60 Hz for each eye was presented using a Windows machine equipped with an OpenGL graphics card (Wildcat series; 3DLABS). Hardware double buffering was used to provide smooth animation. The experimenter's monitor and the video interface of a fiber-optic stimulus presentation system (Silent Vision; Avotec) were driven by the VGA outputs. The field of view was 30 (horizontal) × 23 (vertical) degrees of visual angle, and the focus was fixed at 2 diopters. Binocular presentation was possible through two independently positioned plastic, fiber-optic glasses; however in this study we used monocular stimulation (either left or right eye). The contact lenses for the eyes had matched diopter with an Avotec projector, to focus images on the retina. Positioning was aided by a modified fundus camera (RC250; Carl Zeiss) with an attachment to hold the projector on the same axis of the camera lens. After observing the foveal region, the projector was fixed relative to the animal.

In the awake recordings, the visual stimuli were generated by in-house software written in C/Tcl and used OpenGL implementation. Stimuli were displayed using a dedicated graphics workstation (TDZ 2000; Intergraph Systems, Huntsville, AL, USA) with a resolution of 1,280 × 1,024 and a 60Hz refresh rate. An industrial PC with one Pentium CPU (Advantech) running the QNX real-time operating system (QNX Software Systems) controlled the timing of stimulus presentation, digital pulses to the Neuralynx system (anesthetized) or the Blackrock system (awake), and acquisition of images. Eye movements were captured using an IR camera at 1kHz sampling rate using the software iView (SensoriMotoric Instruments GmBH, Germany). They were monitored online and stored for offline analysis using both the QNX-based acquisition system and the Blackrock data acquisition system.

In the anesthetized recordings, neural activity was recorded in 200 trials of repeated stimulus presentation. Each trial consisted of the same 10s long movie clip, followed by 10s of a blank screen (intertrial). In the awake experiments, two monkeys were trained to fixate on a red square of size 0.2° of visual angle subtended on the eye about 45cm from the monitors, and maintain fixation within a window of 1.5-2° of visual angle. The location of the red fixation square was adjusted to the single eye vergence of each individual monkey. After 300ms of fixation, a moving grating of size 8°, moving at a speed of 12° (monkey H) and 13° (monkey A) per second with a spatial frequency of 0.5 cycles/degree of visual angle and at 100% contrast was presented for 1000ms. The gratings encompassed 8 different directions of motion viz. 0°, 45°, 90°, 135°, 180°, 225°, 270° and 315° (Figure 1B), pseudo-randomized within a block of 8 trials. After 1000ms, a 300ms stimulus-off period preceded the completion of the trial. The monkeys were given a liquid reward (either water or juice) at the end of the trial, if they maintained fixation within the specified fixation window during the entire duration of the trial. Every successful trial was followed by a 1000ms inter-trial period. On average, we found 32% +/-5% of all recorded neurons to be visually modulated. The stimuli, although presented through a stereoscope (due to the data being collected on the same day with other experiments requiring dichoptic viewing conditions), were always presented monocularly in the left eye to remain consistent with the monocular stimulation protocol used in the anesthetized recordings. In both anesthetized and awake recordings, to ensure accurate control of stimulus presentation, a photodiode was attached to the experimenter's monitor permitting the recording of the exact presentation time of every single frame.

In awake recordings spontaneous activity datasets were were collected on days different from those of the task recordings. In anesthetized recordings spontaneous activity datasets were

recorded between periods of visual stimulation. In both awake and anesthetized recordings of spontaneous activity the monitors were turned off. The duration of each spontaneous activity dataset was between 40 to 80 minutes.

Tuning functions and signal correlations

Tuning curves for each detected single unit were computed by averaging the firing rate across trials for each of the 8 presented directions of motion. Signal correlations, defined as the correlation coefficient between the tuning curves of a neuronal pair, were also computed. In addition to classical tuning curves (direction and orientation selectivity), other types of tunings such as inverted tunings, for example, have also been reported in the electrophysiological studies of the macaque prefrontal cortex (Zhou et al., 2012). Because of this variability in the observed tuning properties of detected single units, signal correlation provides a more general measure of response similarity and therefore it was used to investigate the correlation structure that arises from this functional similarity.

Spike count correlations

To compute spike count correlation (r_{sc}) during the anesthetized state, we divided the period of visual stimulation into 10 periods, each being 1000ms long, and considered these periods as different successive stimuli. The intertrial period was also binned in the same way. In the awake data, visual stimulation and intertrial periods were 1000ms long each, thus being consistent with the anesthetized experiments. We estimated spike counts over 1000ms due to the stimulus length used in previous studies of correlated variability. In spontaneous datasets (both anesthetized and awake), the entire length of the recording epoch was split into periods of 1000ms that were

treated as a trial.

The spike count correlation coefficients were computed similar to previous studies in primary visual areas (Bair et al., 2001; Belitski et al., 2008; Ecker et al., 2010). First, for each condition (either presentation of each moving grating in awake experiment or a single bin of movie clip in the anesthetized experiment), we normalized the spike counts across all trials by converting them into *z*-scores (Bair et al., 2001). Then for each pair we computed the Pearson's correlation coefficient for the two vectors z_i and z_j as following:

$$c\,(r_{i,}r_j) \;= E[z_i z_j] \qquad \text{Eq. (1)}$$

After computing $c(r_{i,}r_j)$ for each condition, we averaged across conditions to obtain the correlation value. Equivalently, one can concatenate z for all the conditions in long vectors and find the expectation of their product. To account for possible non-physiological correlations between detected neurons due e.g to electrode shorting, a threshold of 5 standard deviations above the mean correlation value was set and the outliers were discarded.

Quantification of spatial inhomogeneities in correlated variability

We quantified the inhomogeneity in the spatial structure of correlated variability across different conditions and states by computing the mean of the absolute rate (i.e. first differential) of correlation changes across lateral distance. To estimate the first differential with respect to distance we subtracted the mean correlation values of consecutive bins that were significantly different (Wilcoxon rank-sum test, alpha level 0.05). If no significant change between two consecutive bins was observed, the derivative at that point was set to zero.

Author Contributions

Conceptualization, T.I.P.; Methodology, S.S., A.D., V.K. and T.I.P.; Software, S.S., A.D., T.I.P. and J.W.; Formal Analysis, S.S., A.D. and T.I.P.; Investigation, V.K., A.D., T.I.P., S.S. and N.G.H.; Resources, N.K.L. ; Data Curation, A.D., T.I.P., V.K., and S.S.; Writing – Original Draft, T.I.P. , S.S., and A.D.; Writing – Review & Editing: V.K., A.D., T.I.P., N.G.H. and N.K.L.; Visualization, S.S., A.D and T.I.P.; Supervision and Project administration, T.I.P.; Funding acquisition, N.K.L.

Acknowledgments

This study was supported by the Max Planck Society. We thank Britni Crocker and Zeynab Razzaghpanah for help with pre-processing of the data and spike sorting, and Yusuke Murayama and the other technical and animal care staff for excellent technical assistance. We also thank Prof. Andreas Tolias for help with the initial implantations of the Utah arrays and Dr. Michel Besserve and Profs. Christos Constantinidis and Rodrigo Quian Quiroga for their comments on a previous version of this manuscript.

Figure captions

Figure 1: Implantation, visual stimulation and quality of single unit isolation

(A) Location of the implanted array with respect to arcuate and principal sulci and an example of typical waveforms acquired across the implanted cortical patch during a typical recording session in an awake animal.

(B) Anesthetized visual stimulation protocol: 10 seconds of movie clip presentation was interleaved with 10 second long inter-trial (stimulus off) periods for 200 repetitions.

(C) Awake visual stimulation protocol: The macaques initiated each trial by fixating on a red dot for 300ms, following which a drifting sinusoidal grating was presented for 1000ms. After 1000ms of visual stimulation and a 300ms stimulus off period, liquid reward was delivered for successful fixation throughout the trial period. An intertrial period of 1000ms preceded the next trial. Each block of trials comprised eight different motion directions (exemplified by differently coloured arrows) presented in a random order.

(D) Single unit isolation quality: Each column shows the activity recorded from 4 channels recorded in two different datasets, one from each of the two monkeys (monkey H and monkey A). 500 example waveforms for single units (shown as coloured clusters) and noise spikes (multi-unit activity shown as grey clusters) along with the mean waveform in dashed-white, and their corresponding first and second principal components (pc1 and pc2) are shown in the first and second row, respectively. In the last row the first principal component of all the waveforms in a cluster is plotted over time, demonstrating stability of recordings and single unit isolation for periods lasting ~3.5 to 4 hours.

Figure 2: Spatial structure of correlated variability

(A) Spike count correlations (r_{sc}) during visual stimulation (red), intertrial (black) and spontaneous activity (dark blue) as a function of lateral spatial distance (mm) between cell pairs for anesthetized state recordings (error bars represent mean ± SEM).

(B) Same as (A) for awake state recordings.

Figure 3: Distributions of correlated variability across different states and conditions

(A) Distribution of pairwise correlated variability (fraction of pairs) and mean values (dotted lines) during visual stimulation for anesthetized (blue) and awake (pink) recordings. Correlated variability was significantly stronger during anesthesia as a result of a shift in the distribution of pairwise correlations towards positive values.

(B) Same as (A) for anesthetized state recordings during visual stimulation (red), intertrial (black) and spontaneous activity (blue) periods.

(C) Same as (B) for awake state recordings

Figure 4: Quantification of spatial inhomogeneity in the structure of correlated variability

(A) Spatial inhomogeneity in the structure of correlations across different conditions during awake state recordings. Spatial inhomogeneity was quantified by computing the average of the absolute rate of change in the correlation structure across successive distance bins (only those rates significantly different in successive distance bins) (see also Experimental Procedures).

(B) Same as (A) for anesthetized state recordings.

Figure 5: Visual modulation of single unit activity during wakefulness

Spike raster plots and overlaid peristimulus time histograms (PSTHs) for four single units (A-D) across all 8 orientations (0° - 315°). Polar plots for each unit show the preferred direction(s) of motion. The green lines in the centre indicate the resultant length and direction (see Supplemental Experimental Procedures, "Quantification of direction selectivity" for more details). The upper two panels are typical examples of bimodally tuned prefrontal units with significant responses for opposite directions of motion (orientation selective responses). The lower two panels are examples of sharper, unimodal responses for a particular direction of motion.

Figure 6: Effect of functional similarity on the spatial structure of correlated variability

(A) Correlated variability as a function of distance and signal correlation for the pooled data recorded during visual stimulation in the awake state. The colour of each pixel indicates the average correlated variability for pairs that their signal correlation and distance landed in the specific bin. Pixels containing less than 10 pairs are removed (white pixels). Correlated variability values are indicated by the colourbar at the right of the panel. Data were smoothed with a two-dimensional Gaussian (SD of 1 bin) for display purposes.

(B) Correlated variability as a function of distance (similar to Figure 2A and B) among neuronal pairs with signal correlation higher than 0.1 and less than 0.9, i.e. the non-zero upper part of matrix represented in (A) with green line; and among pairs with signal correlation higher than -0.9 and less than -0.1 i.e. the non-zero lower part of matrix represented in (A) with blue line. (Mean ± SEM as error bars)

(C-D) Same as (A) and (B) for the intertrial period. The signal correlation matrix is computed from the visual stimulation period and the correlated variability of these populations in the intertrial period is mapped onto the pixels in C.

(E-G) Three example pairs with high signal correlations and high, low and high correlated variability from the nearest, the intermediate and the farthest distance bins, respectively. The polar plot shows the vector sum of the tuning for each neuron in a given pair while the scatter plots depict their z-score normalized responses. Example raster plots are overlaid with the peristimulus time histograms (PSTH) for the preferred direction of motion. Despite the sparseness in firing for some of the neurons, sharp tunings can be observed (compare raster plots with polar plots).

References

Amir, Y., Harel, M., and Malach, R. (1993). Cortical hierarchy reflected in the organization of intrinsic connections in macaque monkey visual cortex. J Comp Neurol *334*, 19–46.

Angelucci, A., Levitt, J.B., Walton, E.J.S., Hupe, J.-M., Bullier, J., and Lund, J.S. (2002). Circuits for local and global signal integration in primary visual cortex. J Neurosci *22*, 8633–8646.

Bair, W., Zohary, E., and Newsome, W.T. (2001). Correlated firing in macaque visual area MT: time scales and relationship to behavior. J Neurosci *21*, 1676–1697.

Belitski, A., Gretton, A., Magri, C., Murayama, Y., Montemurro, M.A., Logothetis, N.K., and Panzeri, S. (2008). Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. J Neurosci *28*, 5696–5709.

Bullmore, E., and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. Nat Rev Neurosci *10*, 186–198.

Chaudhuri, R., Bernacchia, A., and Wang, X.-J. (2014). A diversity of localized timescales in network activity. Elife *3*, e01239.

Chaudhuri, R., Knoblauch, K., Gariel, M.-A., Kennedy, H., and Wang, X.-J. (2015). A Large-Scale Circuit Mechanism for Hierarchical Dynamical Processing in the Primate Cortex. Neuron *88*, 419–431.

Chelaru, M.I., and Dragoi, V. (2016). Negative correlations in visual cortical networks. Cereb Cortex *26*, 246–256.

Chen, J., Hasson, U., and Honey, C.J. (2015). Processing Timescales as an Organizing Principle for Primate Cortex. Neuron *88*, 244–246.

Cohen, M.R., and Kohn, A. (2011). Measuring and interpreting neuronal correlations. Nat Neurosci *14*, 811–819.

Constantinidis, C., and Goldman-Rakic, P.S. (2002). Correlated discharges among putative pyramidal neurons and interneurons in the primate prefrontal cortex. J Neurophysiol *88*, 3487–

3497.

Constantinidis, C., Franowicz, M.N., and Goldman-Rakic, P.S. (2001). Coding specificity in cortical microcircuits: a multiple-electrode analysis of primate prefrontal cortex. J Neurosci *21*, 3646–3655.

Denman, D.J., and Contreras, D. (2014). The structure of pairwise correlation in mouse primary visual cortex reveals functional organization in the absence of an orientation map. Cereb Cortex *24*, 2707–2720.

Douglas, R.J., and Martin, K.A.C. (2004). Neuronal circuits of the neocortex. Annu Rev Neurosci *27*, 419–451.

Douglas, R.J., and Martin, K.A.C. (2007). Mapping the matrix: the ways of neocortex. Neuron *56*, 226–238.

Douglas, R.J., Martin, K.A.C., and Whitteridge, D. (1989). A Canonical Microcircuit for Neocortex. Neural Comput *1*, 480–488.

Ecker, A.S., Berens, P., Keliris, G.A., Bethge, M., Logothetis, N.K., and Tolias, A.S. (2010). Decorrelated neuronal firing in cortical microcircuits. Science *327*, 584–587.

Ecker, A.S., Berens, P., Cotton, R.J., Subramaniyan, M., Denfield, G.H., Cadwell, C.R., Smirnakis, S.M., Bethge, M., and Tolias, A.S. (2014). State dependence of noise correlations in macaque primary visual cortex. Neuron *82*, 235–248.

Elston, G.N. (2003). Cortex, cognition and the cell: new insights into the pyramidal neuron and prefrontal function. Cereb Cortex *13*, 1124–1138.

Fujita, I., and Fujita, T. (1996). Intrinsic Connections in the macaque inferior temporal cortex. J Comp Neurol *368*, 467–486.

Gochin, P.M., Miller, E.K., Gross, C.G., and Gerstein, G.L. (1991). Functional interactions among neurons in inferior temporal cortex of the awake macaque. Exp Brain Res *84*, 505–516.

Goldman-Rakic, P.S. (1995). Cellular basis of working memory. Neuron *14*, 477–485.

Goldman-Rakic, P.S., and Schwartz, M.L. (1982). Interdigitation of contralateral and ipsilateral columnar projections to frontal association cortex in primates. Science *216*, 755–757.

Hansen, B.J., Chelaru, M.I., and Dragoi, V. (2012). Correlated variability in laminar cortical circuits. Neuron *76*, 590–602.

Harris, K.D., and Mrsic-Flogel, T.D. (2013). Cortical connectivity and sensory coding. Nature *503*, 51–58.

Harris, K.D., Henze, D.A., Csicsvari, J., Hirase, H., and Buzsáki, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. J Neurophysiol *84*, 401–414.

Hazan L., Zugaro, M., and Buzsáki, G. (2006). Klusters, NeuroScope, NDManager: a free software suite for neurophysiological data processing and visualization. J Neurosci Methods *155*, 207–216.

Hubel, D.H., Wiesel, T.N., Stryker, M.P. (1978). Anatomical demonstration of orientation columns in macaque monkey. J Comp Neurol *177,* 361-80

Kadir, S.N., Goodman, D.F.M., and Harris, K.D. (2014). High-dimensional cluster analysis with the masked EM algorithm. Neural Comput *26*, 2379–2394.

Kiani, R., Cueva, C.J., Reppas, J.B., Peixoto, D., Ryu, S.I., and Newsome, W.T. (2015). Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex. Neuron *85*, 1359–1373.

Kritzer, M.F., and Goldman-Rakic, P.S. (1995). Intrinsic circuit organization of the major layers and sublayers of the dorsolateral prefrontal cortex in the rhesus monkey. J Comp Neurol *359*, 131–143.

Latora, V., and Marchiori, M. (2001). Efficient behavior of small-world networks. Phys Rev Lett *87*, 198701.

Latora, V., and Marchiori, M. (2003). Economic small-world behavior in weighted networks. The European Physical Journal B - Condensed Matter *32*, 249–263.

Leavitt, M.L., Pieper, F., Sachs, A., Joober, R., and Martinez-Trujillo, J.C. (2013). Structure of spike count correlations reveals functional interactions between neurons in dorsolateral prefrontal cortex area 8a of behaving primates. PLoS ONE *8*, e61503.

Levitt, J.B., Lewis, D.A., Yoshioka, T., and Lund, J.S. (1993). Topography of pyramidal neuron intrinsic connections in macaque monkey prefrontal cortex (areas 9 and 46). J Comp Neurol *338*, 360–376.

Logothetis, N., Merkle, H., Augath, M., Trinath, T., and Ugurbil, K. (2002). Ultra high-resolution fMRI in monkeys with implanted RF coils. Neuron *35*, 227–242.

Logothetis, N.K., Guggenberger, H., Peled, S., and Pauls, J. (1999). Functional imaging of the monkey brain. Nat Neurosci *2*, 555–562.

Lund, J.S., Yoshioka, T., and Levitt, J.B. (1993). Comparison of intrinsic connectivity in different areas of macaque monkey cerebral cortex. Cereb Cortex *3*, 148–162.

Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature *503*, 78–84.

Martin, K.A.C., Roth, S., and Rusch, E.S. (2014). Superficial layer pyramidal cells communicate heterogeneously between multiple functional domains of cat primary visual cortex. Nat Commun *5*, 5252.

Maynard, E.M., Nordhausen, C.T., and Normann, R.A. (1997). The Utah intracortical Electrode Array: a recording structure for potential brain-computer interfaces. Electroencephalogr Clin Neurophysiol *102*, 228–239.

Melchitzky, D.S., González-Burgos, G., Barrionuevo, G., and Lewis, D.A. (2001). Synaptic targets of the intrinsic axon collaterals of supragranular pyramidal neurons in monkey prefrontal cortex. J Comp Neurol *430*, 209–221.

Miller, K.D. (2016). Canonical computations of cerebral cortex. Curr Opin Neurobiol *37*, 75–84.

Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. Annu Rev Neurosci *24*, 167–202.

Modha, D.S., and Singh, R. (2010). Network architecture of the long-distance pathways in the macaque brain. Proc Natl Acad Sci U S A *107*, 13485–13490.

Murray, J.D., Bernacchia, A., Freedman, D.J., Romo, R., Wallis, J.D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., et al. (2014). A hierarchy of intrinsic timescales across

primate cortex. Nat Neurosci *17*, 1661–1663.

Paradiso, M.A., Meshi, D., Pisarcik, J., and Levine, S. (2012). Eye movements reset visual perception. J Vis *12*, 11.

Pucak, M.L., Levitt, J.B., Lund, J.S., and Lewis, D.A. (1996). Patterns of intrinsic and associational circuitry in monkey prefrontal cortex. J Comp Neurol *376*, 614–630.

Quiroga, R.Q., Nadasdy, Z., and Ben-Shaul, Y. (2004). Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. Neural Comput *16*, 1661–1687.

Rabiner, L.R., McClellan, J.H., and Parks, T.W. (1975). FIR digital filter design techniques using weighted Chebyshev approximation. Proc. IEEE *63*, 595–610.

Rosenbaum, R., Smith, M.A., Kohn, A., Rubin, J.E., and Doiron, B. (2017). The spatial structure of correlated neuronal variability. Nat Neurosci *20*, 107–114.

Rothschild, G., Nelken, I., and Mizrahi, A. (2010). Functional organization and population dynamics in the mouse primary auditory cortex. Nat Neurosci *13*, 353–360.

Sakurai, Y., and Takahashi, S. (2006). Dynamic synchrony of firing in the monkey prefrontal cortex during working-memory tasks. J Neurosci *26*, 10141–10153.

Schulz, D.P.A., Sahani, M., and Carandini, M. (2015). Five key factors determining pairwise correlations in visual cortex. J Neurophysiol *114*, 1022–1033.

Schwartz, M.L., and Goldman-Rakic, P.S. (1984). Callosal and intrahemispheric connectivity of the prefrontal association cortex in rhesus monkey: relation between intraparietal and principal sulcal cortex. J Comp Neurol *226*, 403–420.

Smith, M.A., and Kohn, A. (2008). Spatial and temporal scales of neuronal correlation in primary visual cortex. J Neurosci *28*, 12591–12603.

Smith, M.A., and Sommer, M.A. (2013). Spatial and temporal scales of neuronal correlation in visual area V4. J Neurosci *33*, 5422–5432.

Smith, M.A., Jia, X., Zandvakili, A., and Kohn, A. (2013). Laminar dependence of neuronal correlations in visual cortex. J Neurophysiol *109*, 940–947.

Sompolinsky, H., Yoon, H., Kang, K., and Shamir, M. (2001). Population coding in neuronal

systems with correlated noise. Phys Rev E Stat Nonlin Soft Matter Phys *64*, 051904.

Sussillo, D. (2014). Neural circuits as computational dynamical systems. Curr Opin Neurobiol *25*, 156–163.

Tanigawa, H., Wang, Q., and Fujita, I. (2005). Organization of horizontal axons in the inferior temporal cortex and primary visual cortex of the macaque monkey. Cereb Cortex *15*, 1887–1899.

Timme, N., Ito, S., Myroshnychenko, M., Yeh, F.-C., Hiolski, E., Hottowy, P., and Beggs, J.M. (2014). Multiplex networks of cortical and hippocampal neurons revealed at different timescales. PLoS ONE *9*, e115764.

Tolias, A.S., Ecker, A.S., Siapas, A.G., Hoenselaar, A., Keliris, G.A., and Logothetis, N.K. (2007). Recording chronically from the same neurons in awake, behaving primates. J Neurophysiol *98*, 3780–3790.

Tsujimoto, S., Genovesio, A., and Wise, S.P. (2008). Transient neuronal correlations underlying goal selection and maintenance in prefrontal cortex. Cereb Cortex *18*, 2748–2761.

Ueda, N., Nakano, R., Ghahramani, Z., and Hinton, G.E. (2000). SMEM algorithm for mixture models. Neural Comput *12*, 2109–2128.

Voges, N., Schüz, A., Aertsen, A., and Rotter, S. (2010). A modeler's view on the spatial structure of intrinsic horizontal connectivity in the neocortex. Prog Neurobiol *92*, 277–292.

Yoshioka, T., Blasdel, G.G., Levitt, J.B., and Lund, J.S. (1996). Relation between patterns of intrinsic lateral connectivity, ocular dominance, and cytochrome oxidase-reactive regions in macaque monkey striate cortex. Cereb Cortex *6*, 297–310.

Zhou, X., Katsuki, F., Qi, X.-L., and Constantinidis, C. (2012). Neurons with inverted tuning during the delay periods of working memory tasks in the dorsal prefrontal and posterior parietal cortex. J Neurophysiol *108*, 31–38.

Figure 1

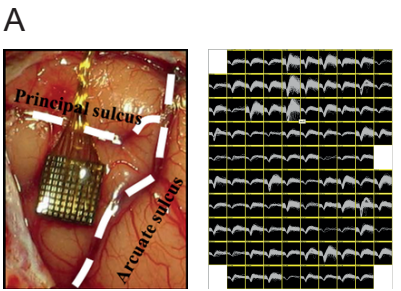
A
Principal sulcus
Arcuate sulcus
B
Movie ON
Movie OFF
trial 1
trial 2
trial 3
trial 200
0
10
20
time (s)
C
1000ms
Intertrial
1000ms
Visual
Stimulation
270°
0°
180°
45°
135°
90°
315°
225°
D
amp
pc1
pc2
pc1
time (s)
'H' Chan 59 SU1
'H' Chan 59 SU2
'H' Chan 94 SU1
'H' Chan 94 SU2
'A' Chan 05 SU1
'A' Chan 09 SU1
'A' Noise Spikes

Figure 2

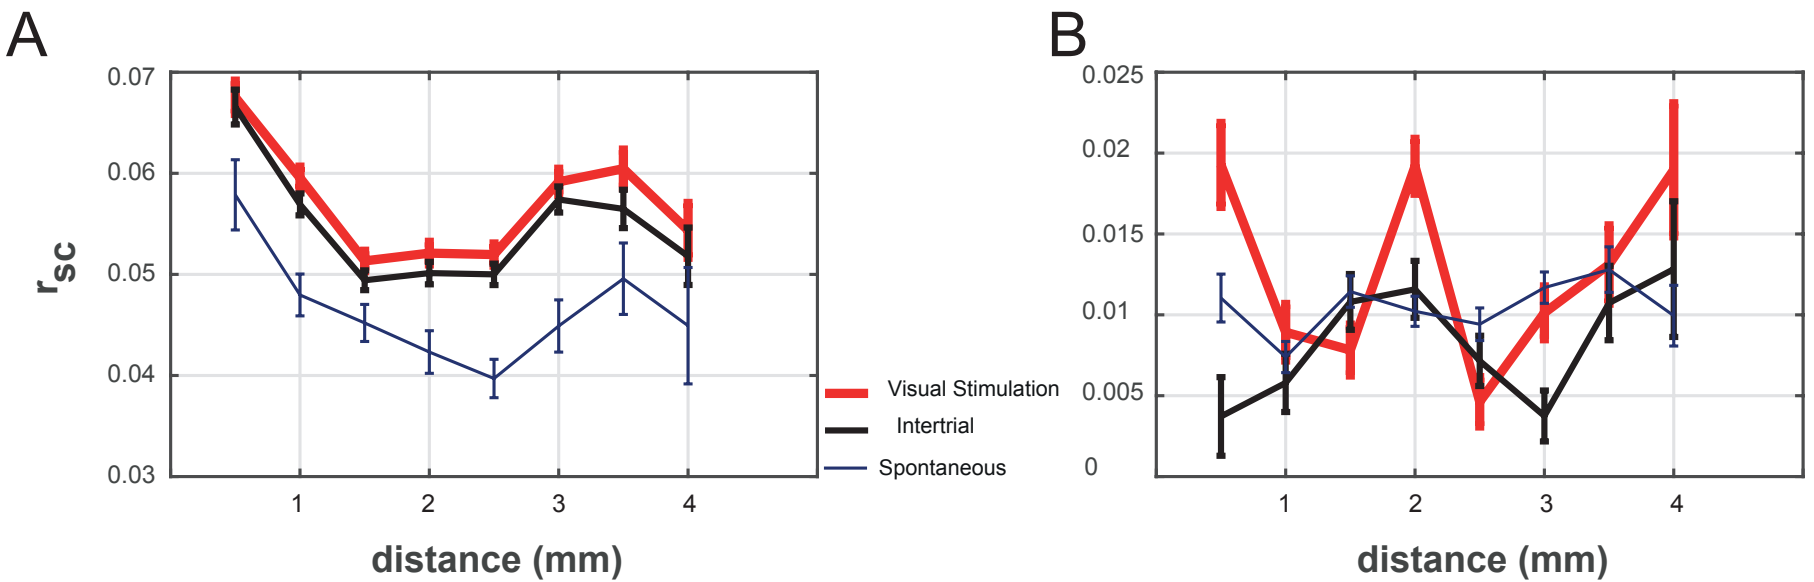
A
B
0.07
0.06
0.05
0.04
0.03
r_{sc}
1
2
3
4
distance (mm)
Visual Stimulation
Intertrial
Spontaneous
0.025
0.02
0.015
0.01
0.005
0
1
2
3
4
distance (mm)

Figure 3

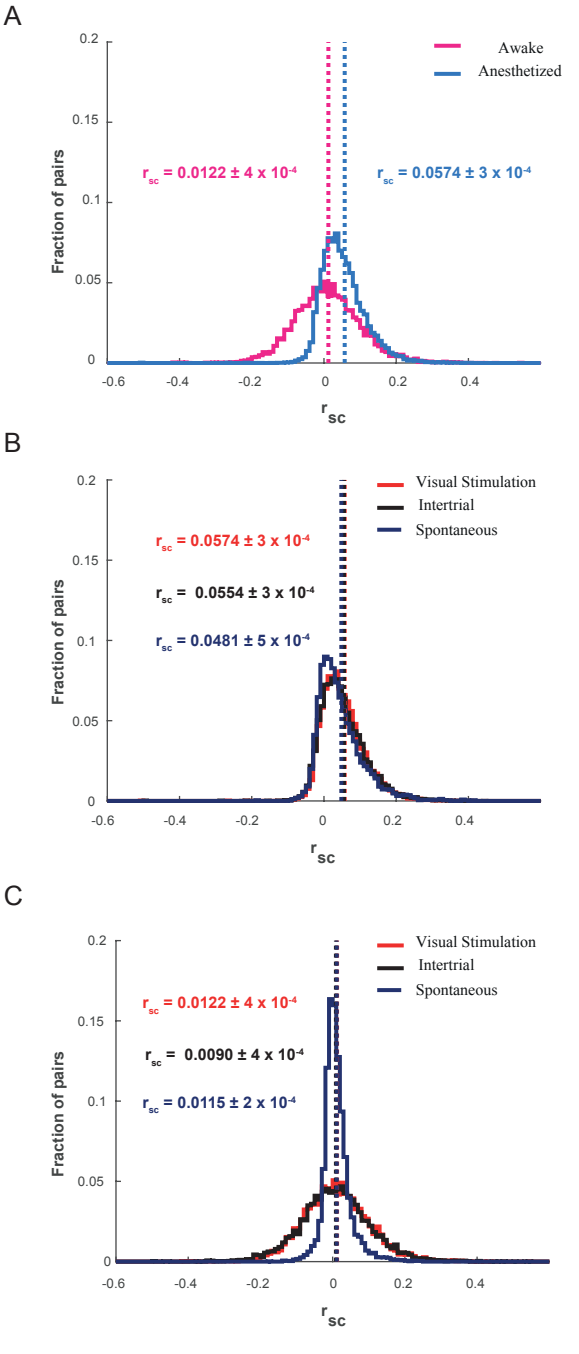
A
Awake
Anesthetized
r_{sc} = 0.0122 ± 4 x 10^{-4}
r_{sc} = 0.0574 ± 3 x 10^{-4}
Fraction of pairs
r_{sc}
B
Visual Stimulation
Intertrial
Spontaneous
r_{sc} = 0.0574 ± 3 x 10^{-4}
r_{sc} = 0.0554 ± 3 x 10^{-4}
r_{sc} = 0.0481 ± 5 x 10^{-4}
Fraction of pairs
r_{sc}
C
Visual Stimulation
Intertrial
Spontaneous
r_{sc} = 0.0122 ± 4 x 10^{-4}
r_{sc} = 0.0090 ± 4 x 10^{-4}
r_{sc} = 0.0115 ± 2 x 10^{-4}
Fraction of pairs
r_{sc}

Figure 4

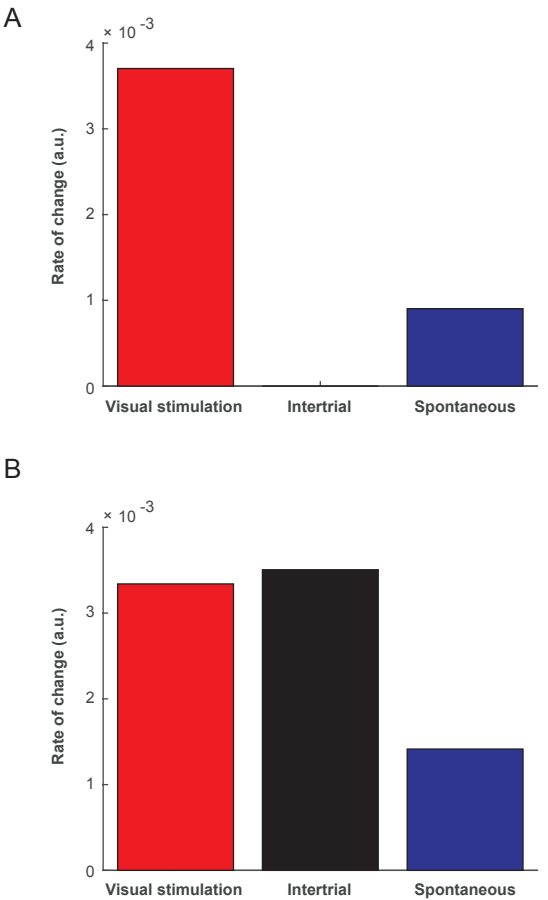
A
× 10 -3
4
3
2
1
0
Rate of change (a.u.)
Visual stimulation
Intertrial
Spontaneous
B
× 10 -3
4
3
2
1
0
Rate of change (a.u.)
Visual stimulation
Intertrial
Spontaneous

Figure 5

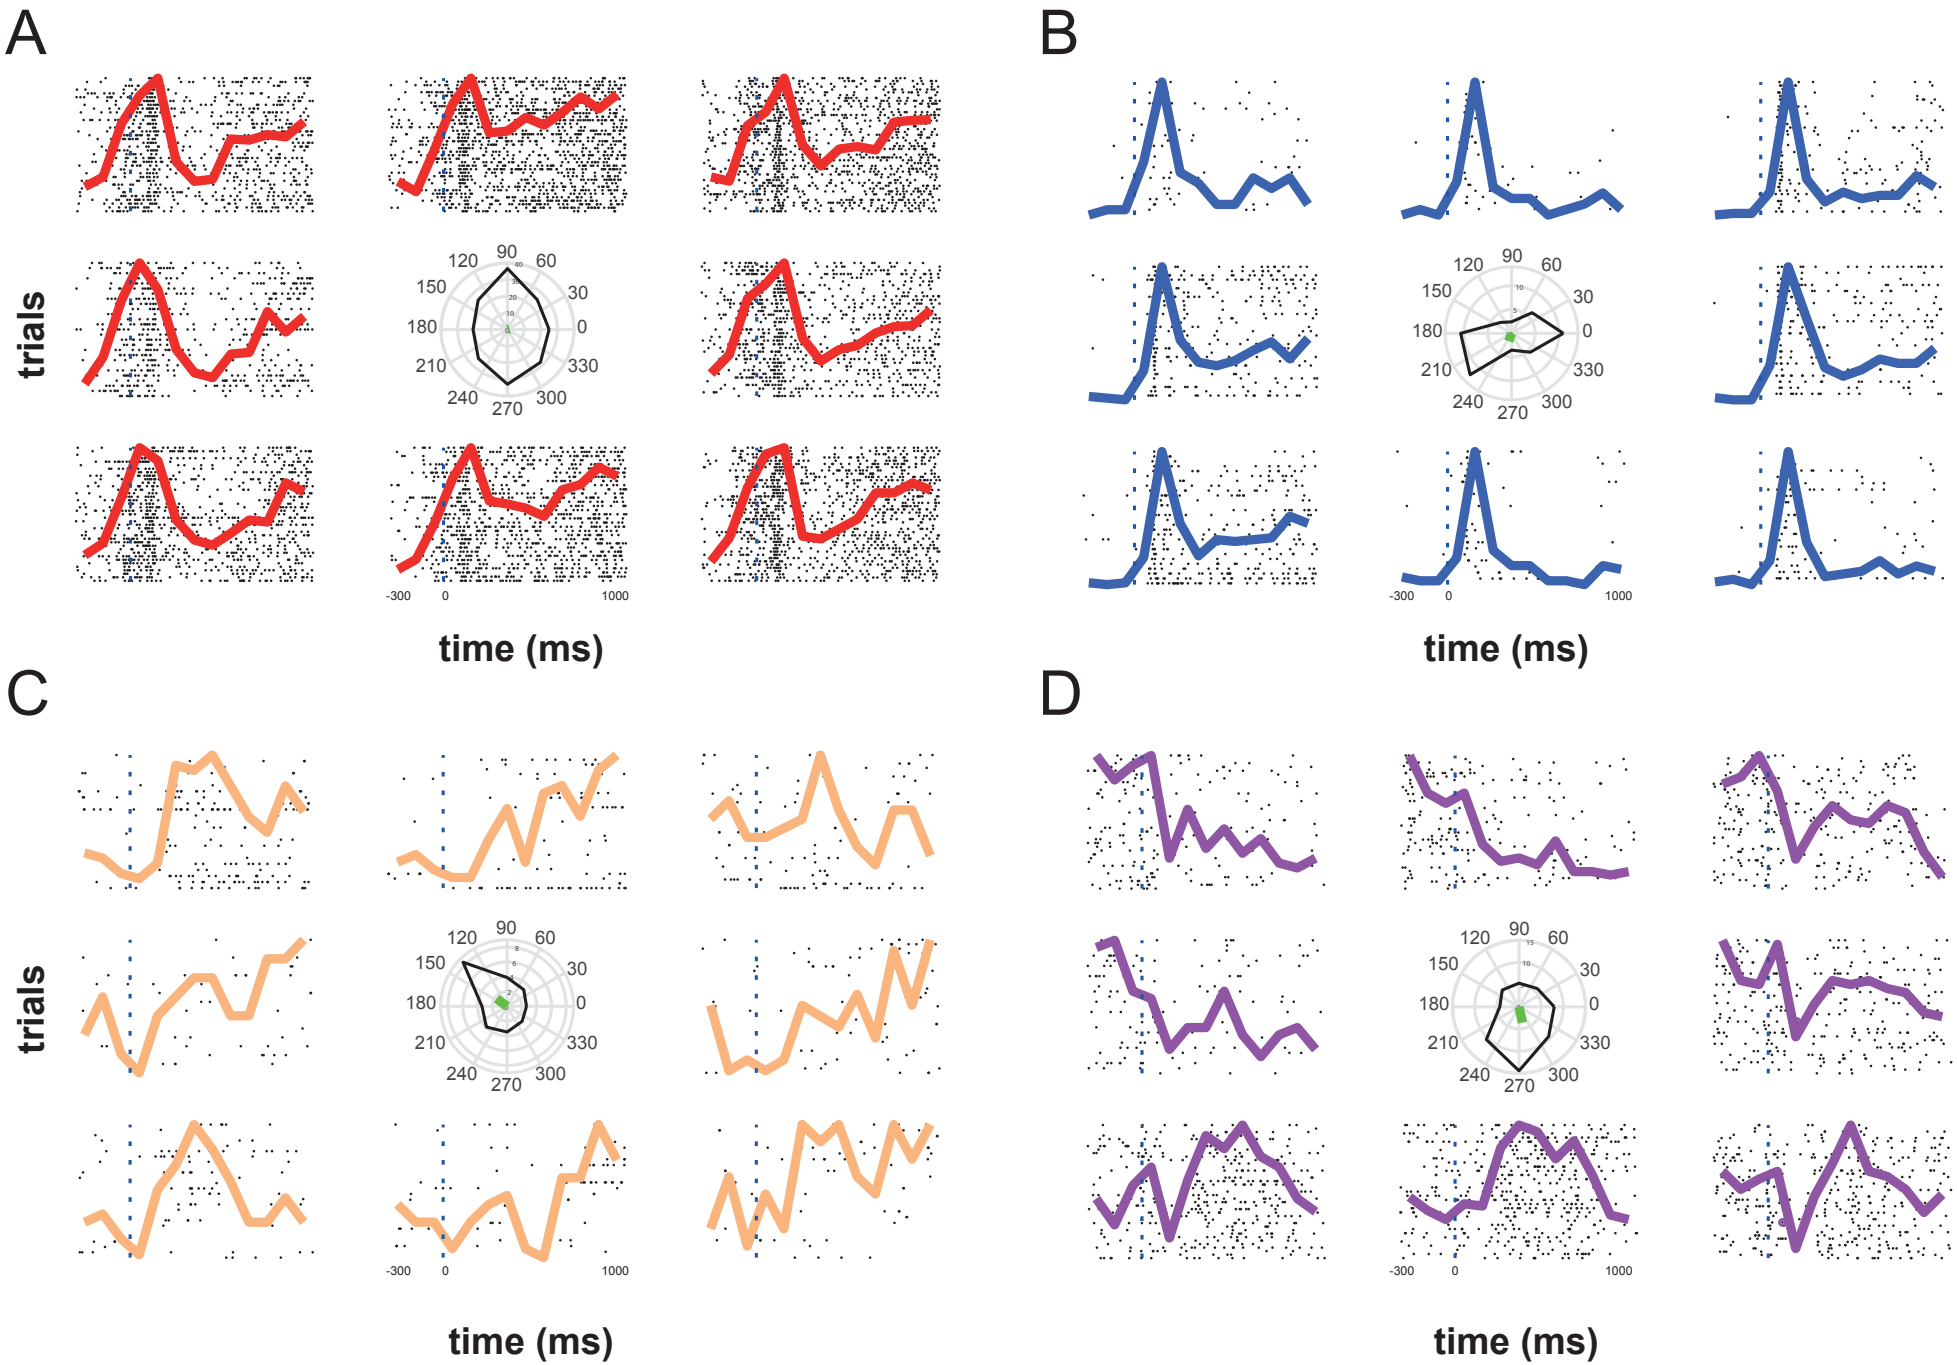
A
B
C
D
trials
trials
90
120
60
150
30
180
0
210
330
240
270
300
-300
0
1000
time (ms)
time (ms)
time (ms)
time (ms)

Figure 6

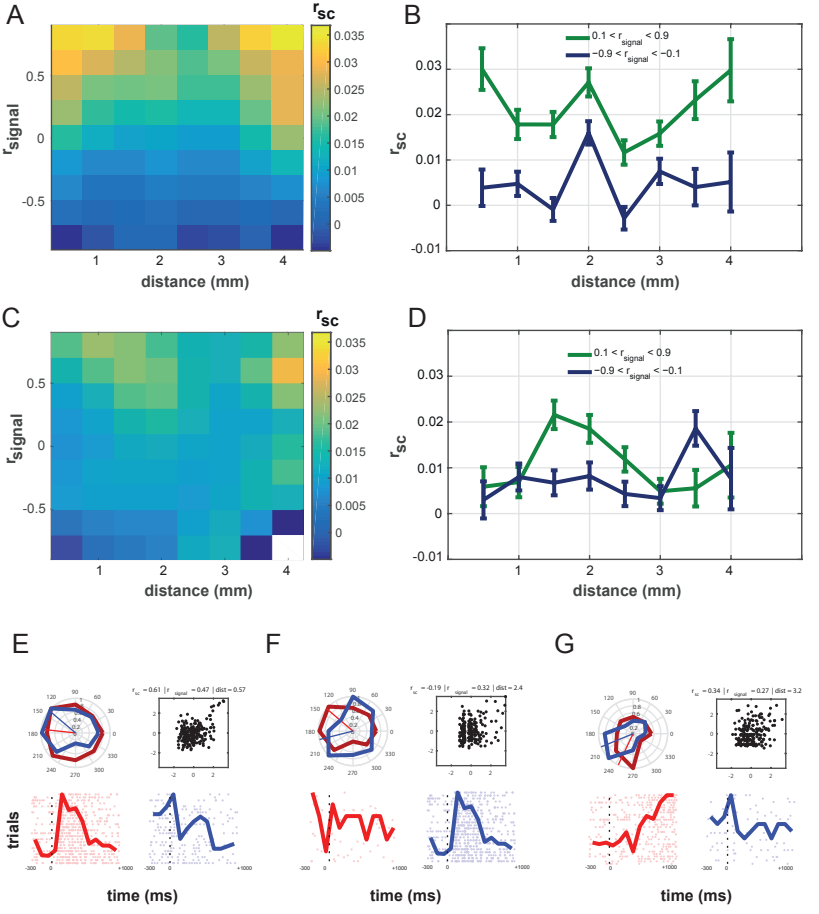
A
B
C
D
E
F
G
r_sc
r_signal
distance (mm)
0.1 < r_signal < 0.9
−0.9 < r_signal < −0.1
trials
time (ms)

Supplemental Experimental Procedures

Quantification of direction selectivity

We quantified direction-selectivity by computing the magnitude of the normalized resultant vector (Constantinidis et al., 2001; Tolias et al., 2005) using CircStat toolbox (Berens, 2009). The directions of motion tested for the selectivity, as mentioned in the main text were equally (45°) spaced from 0 to 360° but not uniformly distributed i.e. there was a small variation in the number of trials for each direction of motion owing to randomisation across the trials. The average spike count across trials belonging to each particular direction was used for computing the magnitude of the normalized resultant vectors (R):

$$R = \left[\frac{\sum_d \overline{r_d} e^{i\theta_d}}{\sum_d \overline{r_d}} \right] \qquad \text{Eq. (1)}$$

Where $\overline{r_d}$ is the average spike count during the presentation of trials with drifting gratings oriented at θ_d degrees.

Surgical preparation and anesthesia

In the anaesthetized experiment, the animals were initially premedicated with glycopyrrolate (0.01 milligrams (mg) per kilogram of body weight (kg), i.m.) and ketamine (15 mg per kg, i.m.). Next, an intravenous catheter was inserted and vital monitors (HP OmniCare/CMS; Hewlett Packard; electrocardiogram, noninvasive blood pressure, CO2, SpO2, temperature) were connected. The monkeys were preoxygenated and anesthesia was induced with fentanyl (3 micrograms (µg) per kg), thiopental (5 mg per kg), and succinylcholine chloride (3 mg per kg) for the intubation of the trachea. The animals were ventilated using an Ohmeda anesthesia machine (Ohmeda), maintaining an end-tidal CO2 of 33 mmHg and oxygen saturation over 95%. Balanced anesthesia was maintained with remifentanil (typical, 1 µg per kg per minute). Mivacurium (5 mg per kg per hour) was used for muscle relaxation. Body temperature was kept constant, and lactated Ringer's solution was given at a rate of 10 milliliters (ml) per kg per hour. During the entire experiment, the vital signs of the monkey and the depth of anesthesia were continuously monitored.

Drops of 1% ophthalmic solution of anticholinergic cyclopentolate hydrochloride were given to each eye to prevent accommodation of the lens and dilation of the pupil. Refractive errors were measured and contact lenses (hard PMMA lenses; Wöhlk) were put on the monkey's eyes with continuous drops of saline throughout the experiment to prevent the eyes from drying. The lenses with the appropriate dioptric power were used to bring the animal's eyes into focus on the stimulus plane.

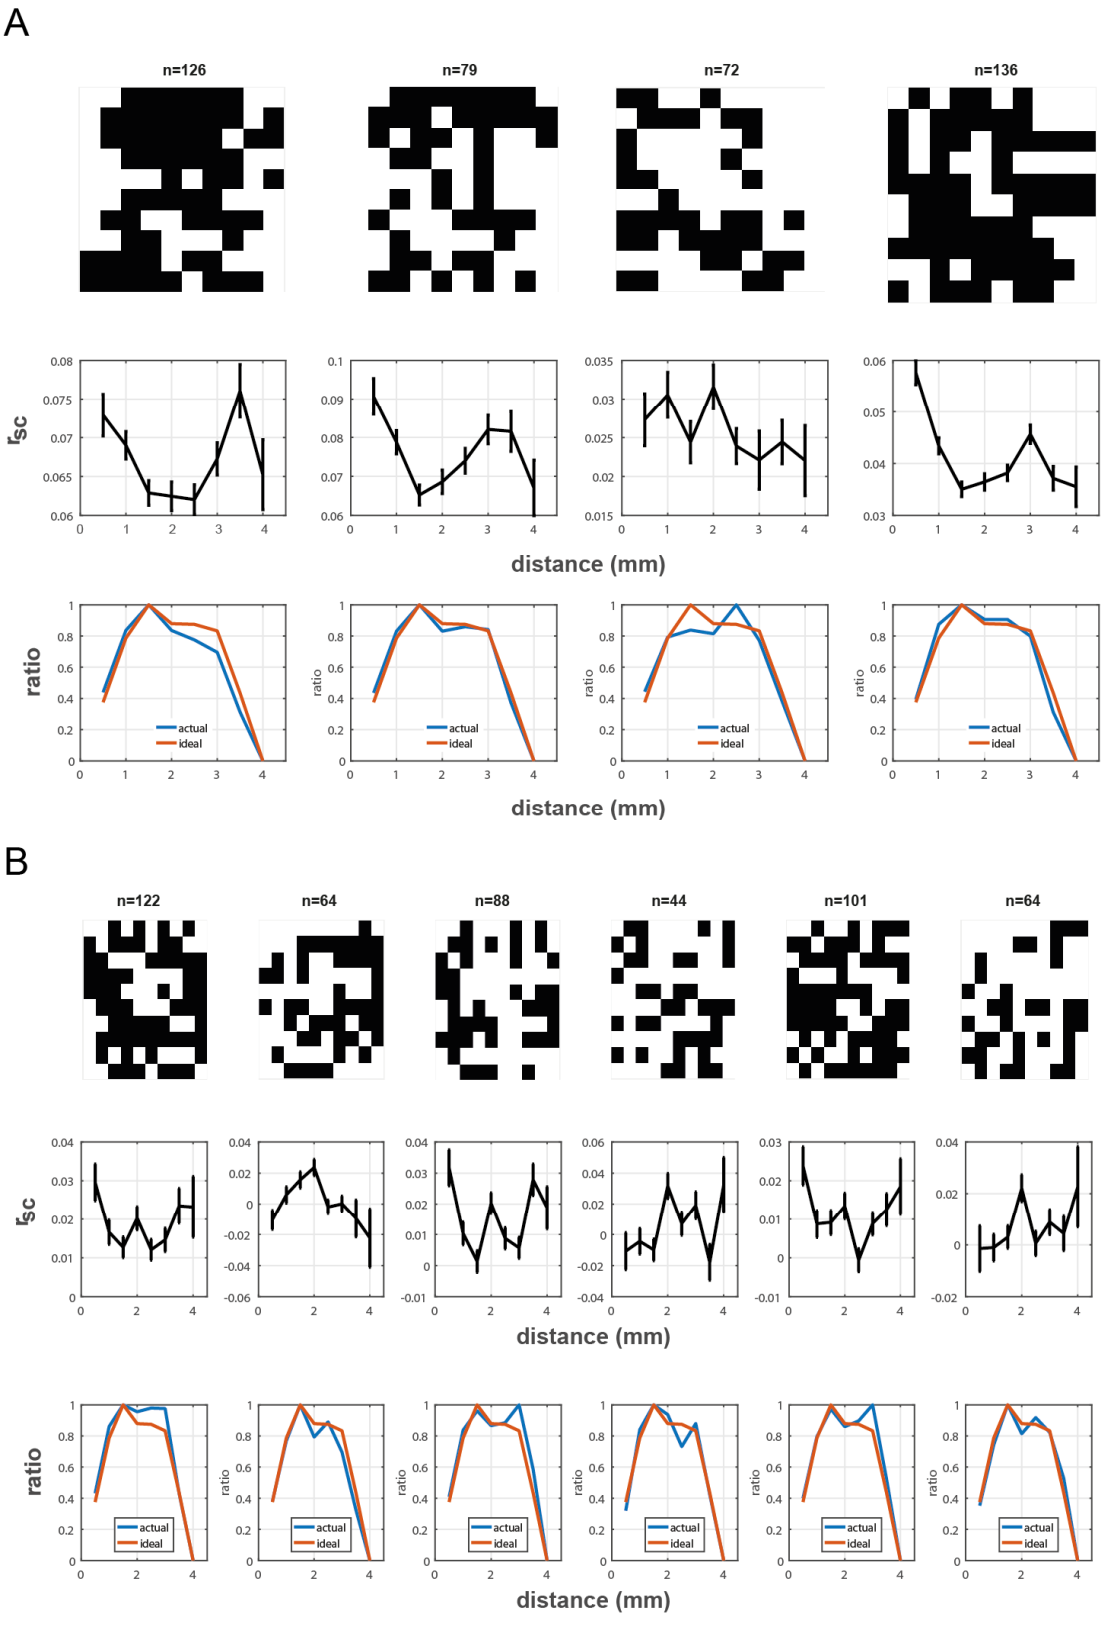

Figure S1: Related to Figure 2: Spatial distribution of single units in anesthetized and awake state recordings

(A) The first row depicts the spatial coverage of the detected single units (*n*) on the Utah array (10 x 10 array of electrodes) for four individual datasets from the anesthetized recordings (each column corresponds to an individual dataset). Black pixels indicate channels where at least one single unit was detected.
In the second row, similar to Figure 2A-B in the main text, correlated variability is depicted as a function of distance during visual stimulation for the same datasets.
In the third row we plot the spatial distributions of the recorded pairs as a function of distance on the Utah arrays (*actual*) for individual datasets. In order to test if the spatial structure of spike count correlations was an artefactual reflection of the spatial distribution of isolated neurons in our recordings, we compared the actual distribution of the pairs with an "ideal" distribution across distance bins. This ideal distribution is derived from a simulated dataset where we could detect equal number of neurons on all sites, indicating the best possible sampling one can achieve with the Utah arrays. To simulate this ideal distribution, we consider a maximum number of detected single neurons on all sites and find the distribution of pairs (aka *ideal* distribution). The largest number of single units among all sites in the recorded dataset is used as the maximum number. For example, if in a given dataset, *n* single units were the largest number of single units across all electrodes, we consider a distribution derived from *n* single units across all sites. The choice of this maximum is not crucial as distributions are normalized. Red and orange traces indicate the actual and ideal distribution of number of pairs (normalized to the peak) in four datasets across all distance bins. The spatial spread of the recorded single units is not significantly different from an ideal distribution.

(B) Similar to figure (A), but for awake state datasets.

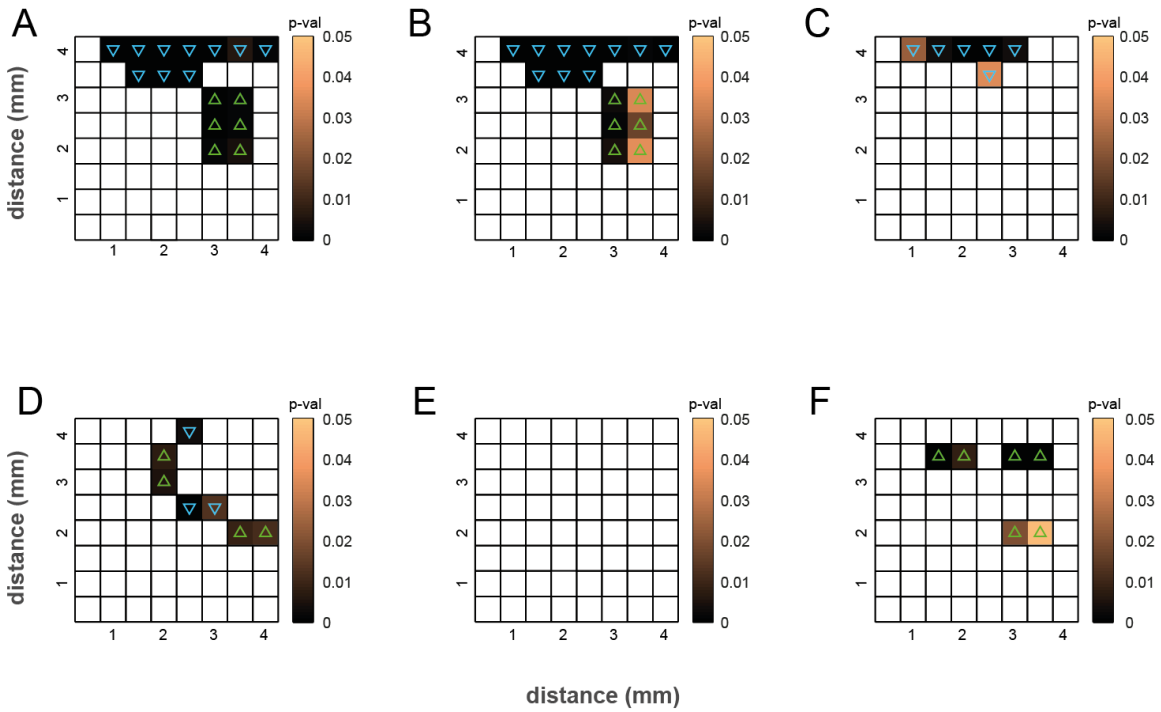

Figure S2: Related to Figure 2: Pairwise comparison of correlations across distance bins

(A-F) Pairwise comparison of correlation distributions (Wilcoxon rank-sum test) across all distance bins for different states and conditions: anesthetized - visual stimulation (A), intertrial (B), spontaneous activity (C) and awake - visual stimulation (D), intertrial (E), spontaneous activity (F). Each plot is an 8 x 8 symmetric matrix (therefore only the upper triangular part is kept) corresponding to all possible pairwise comparisons of 8 distance bins. For example, the p-value of the comparison between 0.5 and 2.5 mm is encoded in the colour of the pixel located in row 1 and column 5 of the matrix. All coloured pixels indicate a significant change since pixels without significant differences ($p>0.05$) were removed (white pixels). Green triangles indicate a significant increase in the correlation values, while blue triangles denote a significant decrease.

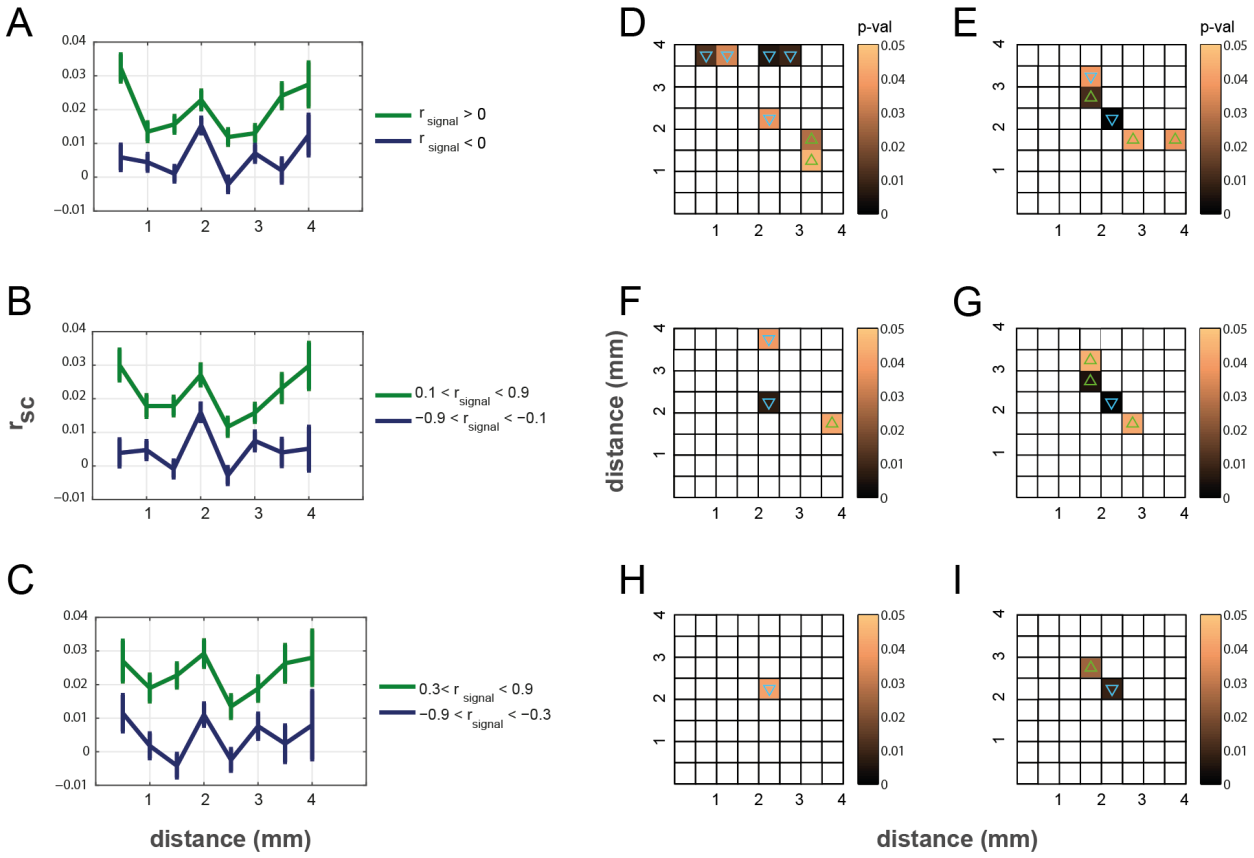

Figure S3: Related to Figure 6: Various thresholds of signal correlation

(A-C) Correlation as a function of distance for different thresholds of signal correlation (similar to figure 6B of the main text): Pairs with positive signal correlations are depicted in green (A, $r_{signal} > 0$; B, $0.1 < r_{signal} < 0.9$; C, $0.3 < r_{signal} < 0.9$) and negative signal correlations are depicted in blue (A, $r_{signal} < 0$; B, $-0.9 < r_{signal} < -0.1$; C, $-0.9 < r_{signal} < -0.3$). (Mean value ± SEM as error bars).

(D, F, H) Similar to the matrix plots in Figure S2, but for different ranges of positive signal correlations (green curves) depicted in (A-C). Each plot depicts the statistical significance (based on the Wilcoxon rank-sum test) of differences in correlations, across all possible pairwise distance bin comparisons.
(E, G, I) Similar to D, F, H but for negative signal correlations (blue curves) depicted in Figure A-C.