

Natural Grouping of Neural Responses Reveals Spatially Segregated Clusters in Prearcuate Cortex

Highlights

- Unsupervised algorithms identify natural functional modules in prearcuate cortex
- Modules are consistent across animals, tasks, and temporal epochs of a task
- Modules segregate more on the basis of “common noise” than task-driven activity
- The “noise” signal is temporally broad-band across several orders of magnitude

Authors

Roozbeh Kiani,
Christopher J. Cueva, ...,
Stephen I. Ryu, William T. Newsome

Correspondence

roozbeh@nyu.edu

In Brief

Kiani et al. apply unsupervised clustering algorithms to multielectrode recordings from monkey cortex, revealing spatially segregated subnetworks that are stable across behavioral contexts. They can be detected from spontaneous activity alone, resembling fMRI “resting state” networks, but at cellular scale.



Natural Grouping of Neural Responses Reveals Spatially Segregated Clusters in Prearcuate Cortex

Roozbeh Kiani,^{1,2,*} Christopher J. Cueva,¹ John B. Reppas,¹ Diogo Peixoto,^{1,4} Stephen I. Ryu,³ and William T. Newsome¹

¹Department of Neurobiology and Howard Hughes Medical Institute, Stanford University, Stanford, CA 94305, USA

²Center for Neural Science, New York University, New York, NY 10003, USA

³Department of Neurosurgery, Palo Alto Medical Foundation, Palo Alto, CA 94301, USA

⁴Champalimaud Neuroscience Programme, Champalimaud Centre for the Unknown, 1400-038 Lisbon, Portugal

*Correspondence: roozbeh@nyu.edu

<http://dx.doi.org/10.1016/j.neuron.2015.02.014>

SUMMARY

A fundamental challenge in studying the frontal lobe is to parcellate this cortex into “natural” functional modules despite the absence of topographic maps, which are so helpful in primary sensory areas. Here we show that unsupervised clustering algorithms, applied to 96-channel array recordings from prearcuate gyrus, reveal spatially segregated subnetworks that remain stable across behavioral contexts. Looking for natural groupings of neurons based on response similarities, we discovered that the recorded area includes at least two spatially segregated subnetworks that differentially represent behavioral choice and reaction time. Importantly, these subnetworks are detectable during different behavioral states and, surprisingly, are defined better by “common noise” than task-evoked responses. Our parcellation process works well on “spontaneous” neural activity, and thus bears strong resemblance to the identification of “resting-state” networks in fMRI data sets. Our results demonstrate a powerful new tool for identifying cortical subnetworks by objective classification of simultaneously recorded electrophysiological activity.

INTRODUCTION

Sensory and motor cortices of the primate brain are often characterized by the presence of topographic maps. For example, primary visual cortex (V1) contains maps of retinotopic space, orientation preference, and ocular dominance (Engel et al., 1994; Katz et al., 1989; LeVay et al., 1975; Van Essen et al., 1984; Wiesel and Hubel, 1974). The boundaries of V1 defined by each of these maps coincide precisely with each other and with architectonic borders as well, reinforcing the notion that V1 is a distinct cortical area with a specific set of functions. Historically, topographies of this nature have been crucial in advancing our understanding of the organization and function of the cerebral cortex (e.g., Felleman and Van Essen, 1991; Hu-

bel and Livingstone, 1987; Mishkin et al., 1983; Zeki and Shipp, 1988).

In contrast, parcellation of the cortex into functional modules is more challenging in association areas where spatial topography may be indistinct or missing altogether. Some areas of the prefrontal cortex (PFC) can be broadly defined by zones of anatomical projections (Carmichael and Price, 1994; Petrides and Pandya, 1999; Preuss, 2007) or general trends in physiological properties. For example, studies in monkeys and humans suggest localization trends within PFC based on sensory input modality (Romanski and Goldman-Rakic, 2002), responses to reward versus punishment (Monosov and Hikosaka, 2012), actual versus hypothetical reward outcomes (Abe and Lee, 2011), and a hierarchy of cognitive control (Badre and D'Esposito, 2009). But outside the frontal eye fields (FEFs) (Bruce et al., 1985), and possibly the frontal lobe “face patches” (O'Scalaidhe et al., 1997; Tsao et al., 2008), sharp boundaries and salient physiological distinctions are rare in PFC. In general, single units recorded in PFC exhibit multiplexed signals of great variety, and neighboring neurons show little evidence of common physiological features that are characteristic of columnar organization in more primary sensory and motor areas.

Here we take a fundamentally different approach to detecting topographic boundaries in PFC. We hypothesized that our limited knowledge concerning topographic organization in frontal cortex may arise from several related limitations of traditional methods for characterizing neuronal activity. First, previous studies have largely relied on a small number of electrodes (usually one), leading investigators to focus on the response properties of individual neurons rather than the population. Second, neural responses are usually characterized by their mean—the first statistical moment of a distribution. Higher moments, especially trial-to-trial fluctuations and response correlations across the population, are frequently not studied, primarily due to lack of simultaneous recordings. And third, neural responses are typically characterized only with respect to task events that are of interest to the experimenter. By breaking these conventional boundaries, it may be possible to discover organizational principles and topographies that have been unknown heretofore.

We approached this problem from a somewhat agnostic perspective. We bypassed some basic limitations of single unit recording by employing multielectrode (Utah) arrays to record

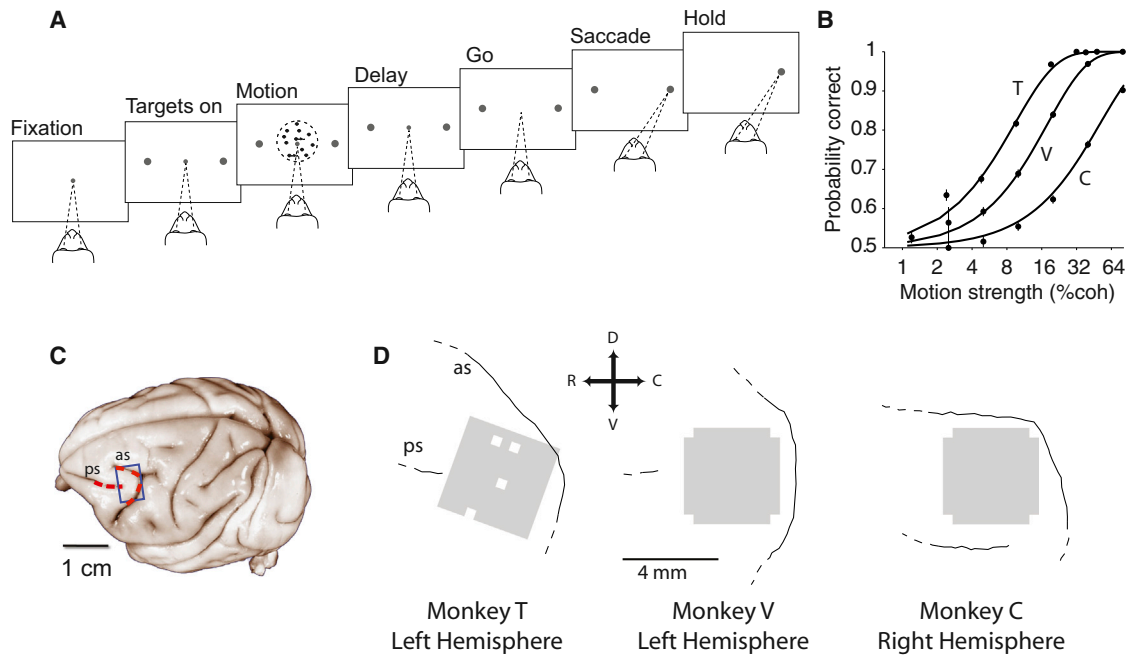


Figure 1. Large-Scale Multielectrode Recordings from the Prearcuate Gyrus during a Direction Discrimination Task

(A) Behavioral task. Monkeys viewed the random dot motion for 800 ms and, after a variable delay, reported the perceived motion direction with a saccadic eye movement. Correct responses were rewarded with juice after a short hold period. The strength and direction of motion varied randomly from trial to trial.

(B) Behavioral performance. The three psychometric functions depict performance for the three monkeys (T, V, and C), averaged across all sessions. Psychophysical thresholds were 9.3% coherence for monkey T, 17.9% coherence for monkey V, and 51% coherence for monkey C. Monkey C's perceptual sensitivity was poor relative to most animals; threshold remained high despite months of training. The results in this paper, however, do not depend upon perceptual sensitivity. Our only requirement is that the animal was under behavioral control during task performance, which is demonstrated by the regular psychometric function.

(C) Target area (blue box) for implantation of the multichannel electrode array on the prearcuate gyrus. Arcuate (as) and principal (ps) sulci are marked with red dashed lines on the surface of a typical macaque brain (University of Wisconsin Brain Collection).

(D) The actual location of each array with respect to arcuate and principal sulci. The white squares show the ground pins. In monkey C, the array could not be placed at the concavity of arcuate sulcus due to the unusually short distance between the arcuate and the posterior termination of the principal sulcus. Dashed lines at the end of a sulcus indicate the sulcus extends in this direction beyond our craniotomy.

simultaneously from tens to hundreds of units at regularly spaced intervals across a specific region of PFC. Second, we used unsupervised algorithms to identify natural groupings of neurons based on their response covariation, both task driven and task independent. Finally, we projected the objectively identified groupings of neurons back onto the arrays to determine whether they were spatially segregated in a topographic manner.

We report recordings from the prearcuate gyrus, a region of PFC that carries visual, cognitive, and eye movement-related signals in a variety of behavioral tasks (Constantinidis and Goldman-Rakic, 2002; Hussar and Pasternak, 2009; Kiani et al., 2014; Kim and Shadlen, 1999; Lennert and Martinez-Trujillo, 2013; Mante et al., 2013). The prearcuate gyrus is traditionally divided into the “core” FEF, located in the rostral bank and lip of the arcuate sulcus, and area 8Ar, located between the arcuate sulcus and the posterior tip of principal sulcus (Gerbella et al., 2007; Schall, 1997; Stanton et al., 1989). Area 8Ar offers a convenient target for dense multielectrode arrays because it is relatively flat. It is unknown if area 8Ar is a homogenous piece of cortex or divides further into smaller subregions. Moreover, electrophysiological recordings are generally considered insufficient to detect the boundary between FEF and 8Ar or to explore sub-

divisions of area 8Ar, because the neurons appear to have similar response properties across the prearcuate gyrus (Constantinidis and Goldman-Rakic, 2002; Hussar and Pasternak, 2010; Kim and Shadlen, 1999).

Here we show that the recorded population in area 8Ar is not homogenous and can be divided into smaller subnetworks based on task-independent covariation of neural responses. The subnetworks are spatially segregated within the prearcuate gyrus, revealing a topography that is defined at the population level by measurements of large-scale, simultaneous recordings. The prearcuate subnetworks may reflect novel areal boundaries within area 8Ar or pronounced interanimal variation of known boundaries (see Discussion). Our new approach will be valuable for detecting boundaries of both kinds as large-scale array and optical recordings become increasingly common in the future.

RESULTS

We used 96-channel multielectrode arrays to record from neural populations in area 8Ar of the prearcuate gyrus (Figure 1) while our subjects, three macaque monkeys, performed a direction discrimination task (Britten et al., 1992; Kiani et al., 2008). On

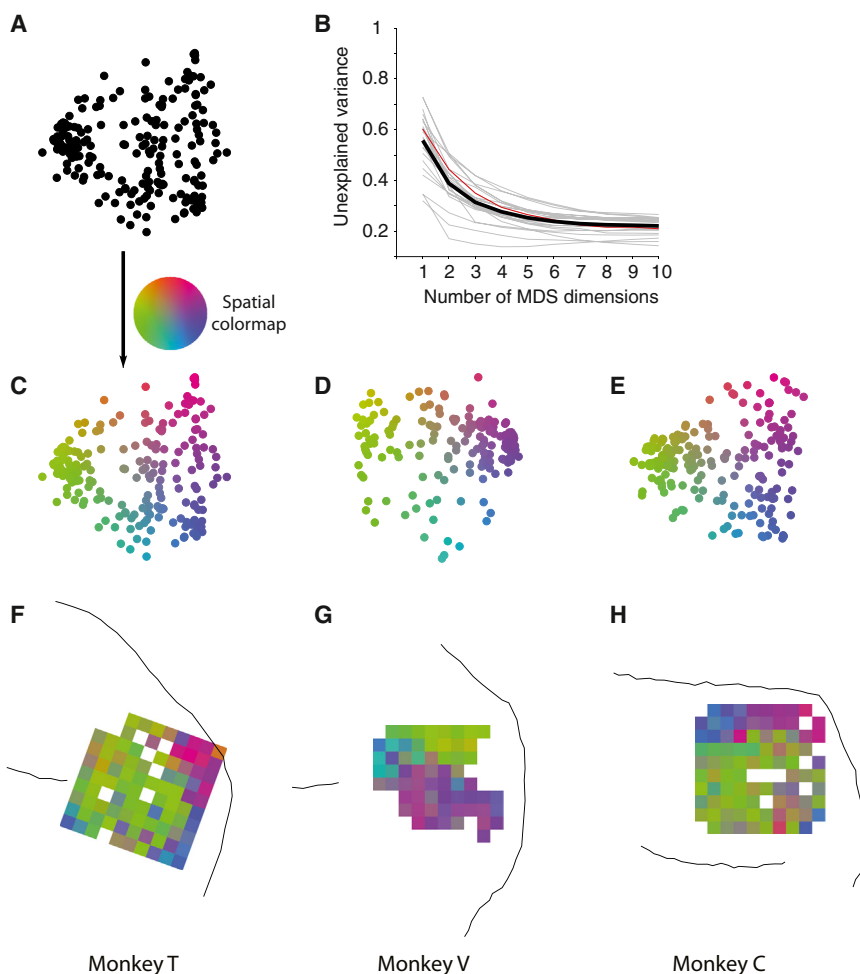


Figure 2. Spatial Topography in Prearcuate Gyrus

(A) 2D depiction of recorded units based on response correlations in an example session. In this depiction, each point represents one unit, and the Euclidean distances between units represent the dissimilarity of their responses ($1 - \text{correlation coefficient}$) across the session. Isomap MDS was used to create this map.

(B) Unexplained variance as a function of the number of MDS dimensions suggests that the dissimilarity matrix is low dimensional. Two dimensions capture a large fraction of variance across sessions (mean = 61.2%). Gray lines represent individual sessions. The thick black line is the average. The red line represents the example session in (A).

(C) The units of the example session in (A) are colored according to a 2D color map in which hue represents radial angle and saturation represents eccentricity.

(D and E) 2D depictions of example sessions in the other monkeys.

(F–H) Projection of the unit colors onto the recording electrodes reveals spatial topography (clustering of colors) within the recording area. White squares correspond to ground pins or to electrodes that failed to record a unit in the depicted session. If an electrode recorded from more than one unit, the average color of the units is projected onto that electrode.

each trial the monkey viewed a patch of randomly moving dots for 800 ms. After a delay period of variable length, the monkey reported the perceived motion direction by making a saccadic eye movement to one of the two available targets. All monkeys were trained on the task before implantation of the recording arrays and showed stable performance throughout the experiments.

The multielectrode array covered 4×4 mm of the cortical surface and enabled us to record simultaneously from multiple single- and multineuron units in a significant portion of the prearcuate gyrus. Consistent with previous studies (Constantinidis and Goldman-Rakic, 2002; Hussar and Pasternak, 2009; Kiani et al., 2014; Kim and Shadlen, 1999; Lennert and Martinez-Trujillo, 2013; Mante et al., 2013; Robinson and Fuchs, 1969), we observed a variety of response properties in different epochs of the direction discrimination task, including visual, decision-related, and perisaccadic signals (see Figure S1 available online).

To explore the presence of functionally specialized circuits within the population of recorded units, we searched for natural groupings of neurons based on temporal covariation of activity over the entire course of an experiment. To do so, we first quantified the responses of each unit as a vector of time-varying firing rates in 30 ms bins from the beginning to the end of the recording session as described in Experimental Procedures (note that we

ignore task events in this first-pass analysis). Then we measured the dissimilarity of physiological activity for each possible pair of units as one minus the Pearson's correlation (r) of their firing rate vectors

(Equation 2). Thus our dissimilarity index ($1-r$) varies between 0 and 2 for perfectly correlated and perfectly anticorrelated units, respectively. We will use the term “dissimilarity matrix” to refer to the set of dissimilarity indices for all possible pairs of units on a given array. Compatible with previous studies in visual cortex (Smith and Kohn, 2008) and PFC (Constantinidis and Goldman-Rakic, 2002; Leavitt et al., 2013), response dissimilarity increased with distance between recording electrodes and decreased with the duration of the spike count measurement window (Figure S2).

Having calculated the dissimilarity matrix for a given experiment, we visualized the dissimilarities between units using multi-dimensional scaling (MDS). Figure 2A shows a MDS map for an example session from monkey T. Each point in the map represents one unit, and the Euclidean distance between any two units represents the pairwise response dissimilarity of those two units—as well as possible for a 2D projection. Thus, neighboring units in a map are more strongly correlated than distant units. Figures 2C–2E illustrate one MDS map from each monkey, with each unit colored according to a 2D color map that will be used in subsequent analyses. Plotting the unexplained variance of the dissimilarity matrix as a function of the number of MDS dimensions showed that four dimensions were often adequate to

explain the dissimilarity matrix (Figure 2B). The 2D projections captured a large portion of the variance (61.2%) and are highly informative about the structure of the dissimilarity matrix.

The distribution of units in the MDS maps does not appear homogeneous, suggesting that the units can be divided into physiologically distinct clusters (SigClust, $p < 10^{-8}$ for each illustrated experiment). We will refer to these clusters as subnetworks, or “subnets” for short. For 14 of the 25 recording sessions (56%), the neural data clustered into two statistically distinct subnets (SigClust, $p < 0.05$). The remaining 11 sessions showed the same trends, even though they were not individually significant (MDS maps not shown, but see spatial maps for all sessions, Figure S3).

Interestingly, the subnets appear to occupy distinct regions of the recording array (Figures 2F–2H). Note that the physical location of the units on the array played no role in our calculation of the dissimilarity index; note also that a spatial map cannot be directly inferred from the mere presence of clusters in Figures 2C–2E. Projecting the units back onto the recording arrays, however, reveals that the units that clustered together in the MDS maps also tended to form spatial clusters on the array. Furthermore, the locations of the clusters were highly consistent from one session to another in each monkey (Figure S3), revealing a characteristic spatial topography based on response covariation among prearcuate units.

Subnet Identity Is Driven by Correlated Noise, Not by Task-Related Activity

A potential explanation of the subnets is that they are driven solely by task-related events. The response vectors used for our first-pass calculation of dissimilarity indices are a continuous function of time across the entire experiment, including all trial epochs as well as the intertrial interval. Thus units that respond strongly to onset of the targets might cluster in one subnet, while units that are more active before and after saccades might form a different subnet. Even if we restrict our analysis to a particular trial epoch, correlated responsiveness (and thus subnet clusters) might still emerge due to differential tuning of units to visual motion direction or to saccade direction and amplitude. We therefore conducted two further analyses to explore the effect of task-related events on our subnet classifications.

In the first analysis, we recalculated response dissimilarities as described above, restricting the analysis to single time epochs during the trial. To analyze the motion-viewing interval, for example, we created a new response vector for each unit by concatenating responses from the motion-viewing interval across the entire experiment, excising all other intervals. We performed this analysis separately for six nonoverlapping time epochs (see [Experimental Procedures](#)) which capture different aspects of neural activity in the direction discrimination task, including responsiveness to visual target onset, motion stimulus, decision formation, pre- and post-saccadic activity, and spontaneous activity during the intertrial interval. Despite these differences, the structure evident in the MDS plots was consistent across all epochs, including the intertrial interval (example experiment, Figures 3A–3F) and was consistent with the dissimilarity structure calculated for the entire session (Figure 2C). Conse-

quently, the spatial topography on the arrays was also replicated independently for each epoch (Figure S4). These data suggest that subnet clustering does not derive primarily from task-related signals such as visual input, motor output, decision-making, or motor planning.

To quantitatively assess the consistency of response dissimilarities across trial epochs, we calculated an ‘alignment score’, which is simply the correlation coefficient of the dissimilarity matrix for a particular epoch with the dissimilarity matrix for the entire session. Figure 3G shows that the alignment scores were consistently high for all temporal epochs, confirming the impression gleaned from visual inspection of the MDS plots. The high alignment scores did not result simply from overlap of data for the individual epochs with whole-session data. We obtained similarly high alignment scores from dissimilarity matrices calculated for nonoverlapping trial epochs (data not shown). Also, the consistency of response dissimilarities was not due simply to a lack of neural responses. Visual target onset, motion viewing, saccade preparation, and execution significantly modulated the activity of the prearcuate population (Constantinidis and Goldman-Rakic, 2002; Hussar and Pasternak, 2010; Kim and Shadlen, 1999). Importantly, the analyses in Figure 3 confirm that task-related changes in activity exert little effect on the subnets defined by the structure in response covariation.

In the second analysis, we further explored task-driven effects by breaking the responses of individual neurons into task-evoked and residual components. The task-evoked responses for individual neurons are defined as the expectation of response magnitude (average response) for each unique combination of motion strength, motion direction, and saccade, which are the task parameters that are controlled or monitored by the experimenter. The residual responses of individual neurons are the trial-to-trial fluctuations around the corresponding means. In essence, we composed two new response vectors for each unit in each experiment, one composed of the response expectation (the average response) in each time bin with trial-to-trial variability removed, and the other composed only of the trial-to-trial residual activity following subtraction of the mean from each time bin. We then recalculated response dissimilarities for each component—task-evoked and residuals.

The response dissimilarities based on residuals align closely with the whole-session response dissimilarities, both qualitatively for individual experiments (Figure 4A, upper row—recall that the color of each data point is maintained from the whole-session analysis) and quantitatively across all experiments (Figure 4B). Using the residual responses of individual epochs establishes statistically significant subnets in 9–12 sessions (36%–48%), depending on task epoch. The slight drop compared to whole-session maps is due to the reduced data available for the analysis. Because the MDS maps for all temporal epochs are well aligned to the whole-session maps, the spatial topography obtained by projecting units back onto the array is maintained as well (Figure S5, upper row). In contrast, response dissimilarities based on task-evoked components (the expectation) are poorly aligned with whole-session data (Figure 4A, lower row; Figure 4C), although the alignment scores are significantly above zero, $p < 0.05$. Unsurprisingly given the MDS plots,

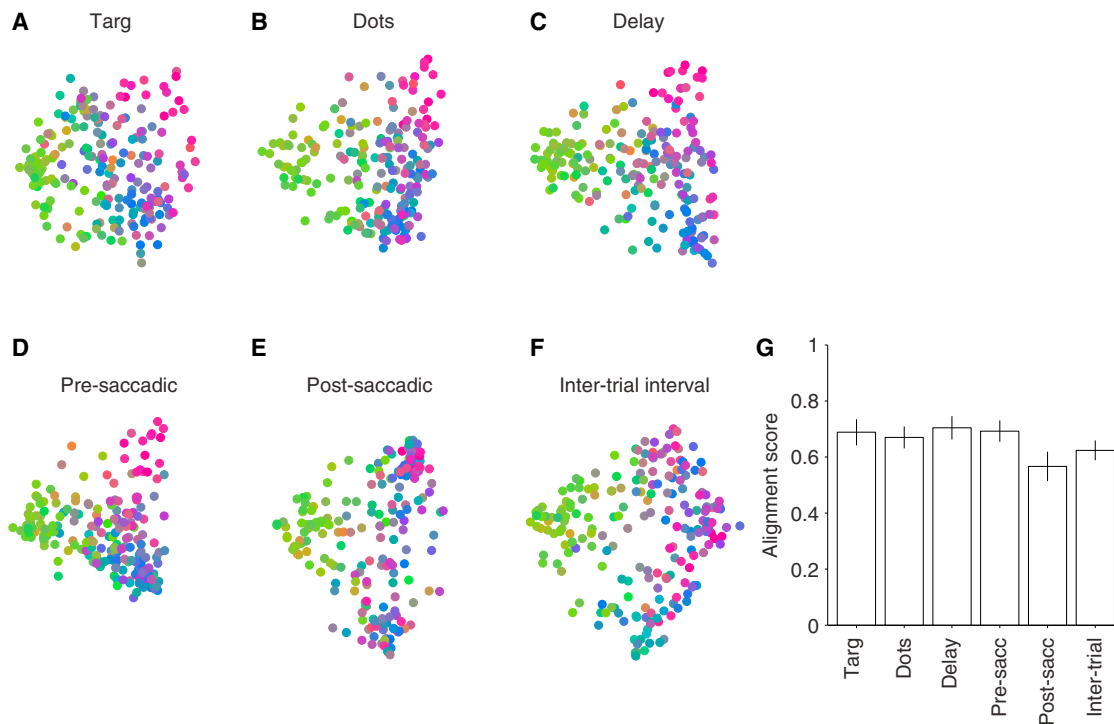


Figure 3. Topography in the MDS Plots, and Thus Spatial Topography on the Arrays, Is Stable across Task Epochs

Same experiment as in Figure 2C.

(A–F) MDS plots calculated independently for six temporal epochs in the task (see [Experimental Procedures](#)). Each unit inherited the same color assigned to it in the whole-session MDS map in Figure 2C. Thus, clustering of units with similar colors indicates that the observed topography is preserved across task epochs. (G) To quantify the preservation of topography, we calculated the correlation of the whole-session dissimilarity matrix with epoch-based dissimilarity matrices (alignment score). The bars show the average alignment scores across sessions. Error bars represent 95% confidence intervals.

task-evoked average responses also fail to fully replicate the spatial topography on the array (Figure S5, lower row).

Thus, it is highly unlikely that task-evoked responses in individual neurons underlie the existence of spatially topographic subnets. Rather, the subnets exist mainly because of what is commonly termed “correlated noise” in traditional electrophysiology experiments. The consistency of the noise structure suggests that it can be informative about network connectivity (Kohn et al., 2009; Ringach, 2009; Tsodyks et al., 1999), even though the functional benefit is not obvious. We revisit this issue in the [Discussion](#).

The Signal that Underlies Subnets Is Temporally Broadband

To better characterize the nature of the residual noise signals that underlie the subnets, we recalculated the dissimilarities of the residual signals within nine temporal frequency bands, from 0.01 Hz to 16.7 Hz (see [Supplemental Information](#)), and measured their alignment to whole-session response dissimilarities. MDS plots for all temporal frequency bands exhibit clustering that is similar to whole-session clustering (Figure S6A), an impression that is confirmed quantitatively by the alignment scores (Figure S6B). Although the best alignment with whole-session data was obtained for 1–4 Hz (roughly delta-band), the alignments were generally good across all frequencies.

Subnets and Response Dissimilarities Are Stable across Different Tasks

Because task-related effects were minimal in the analyses presented above, we hypothesized that the subnets for a particular array would be stable across behavioral tasks. We tested this hypothesis by analyzing data obtained from the same arrays in a visually-guided, delayed saccade task (Figure 5A; see [Experimental Procedures](#)). This task differed from the direction discrimination task in several ways: only one target was presented on the screen on each trial, the location of the target varied substantially from trial to trial, the monkey never viewed the random dot stimulus, and at the time of saccade there was no uncertainty about reward. Figure 5B depicts the MDS plot and spatial topography map for an example experiment in monkey T, which are qualitatively similar to equivalent data from the direction discrimination task (Figures 2C and 2F). Figure 5C shows quantitatively that the dissimilarity matrices calculated from the delayed saccade task are highly aligned with those obtained from the direction discrimination task, across multiple experiments in each animal (Mantel’s test, monkey T, $r = 0.73$, $p < 0.001$; monkey V, $r = 0.54$, $p < 0.001$; monkey C, $r = 0.62$, $p < 0.001$). Moreover, the dissimilarities in the delayed saccade task, like those in the discrimination task, were driven largely by correlated noise (residuals) as opposed to task-evoked responses (Figures S8B–S8D).

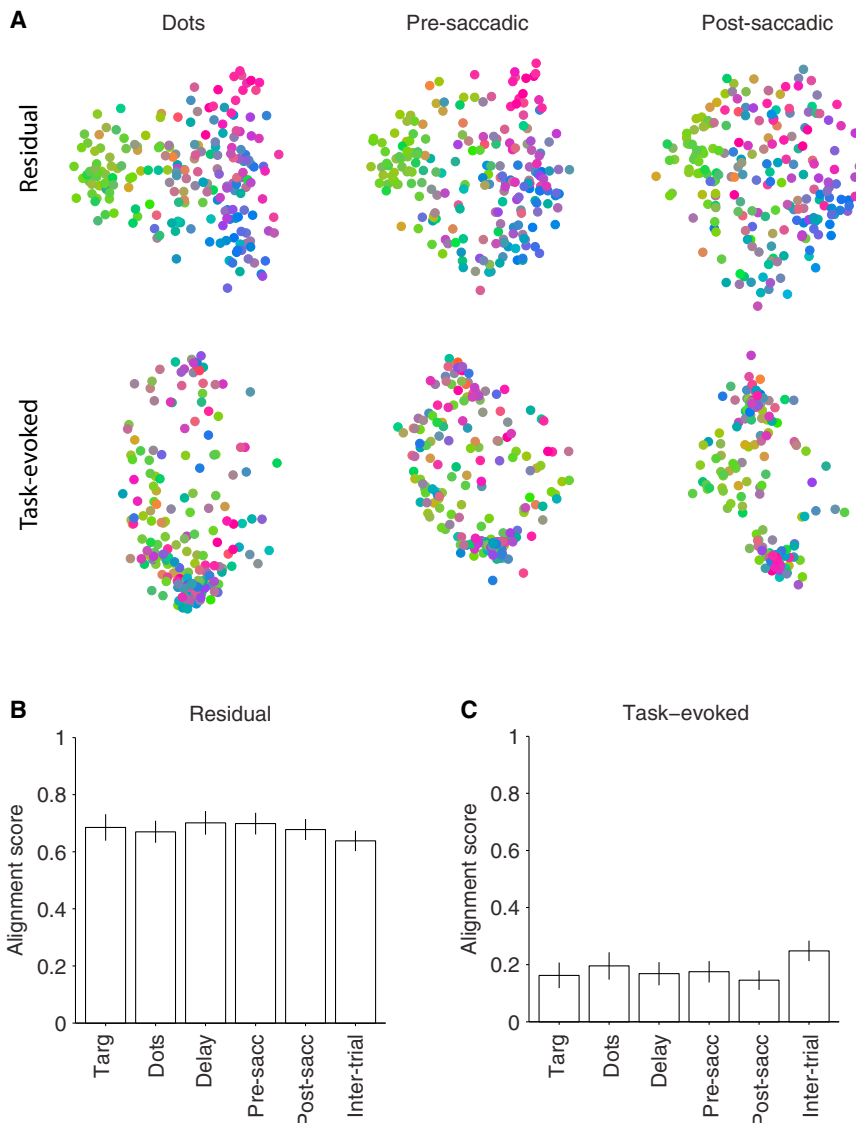


Figure 4. Common Noise Is the Main Underlying Factor for the Topography

(A) 2D plots of units based on task-evoked and residual responses for the example session in Figure 2C. The measured neural responses in each trial epoch consisted of a task-evoked component (the mean across trials with similar motion direction, motion strength, and choice) and a residual component (the variation around the mean). We recomputed dissimilarities for all six temporal epochs based on the task-evoked and residual components. MDS plots are shown for three epochs. The unit colors are inherited from Figure 2C. MDS maps are largely preserved for residual responses, but not for task-evoked responses.

(B) Alignment scores of dissimilarity matrices for the residual responses with those for the whole-session responses in six temporal epochs. The bars show average alignments across sessions.

(C) Alignment scores of dissimilarity matrices for the task-evoked responses with those for whole-session response across the same sessions. Error bars represent 95% confidence intervals.

and were substantially independent of task epoch (Figure S8B, upper row).

The Subnets Exhibit Different Physiological Properties

The existence of spatially segregated neural clusters in our study raises the possibility that neurons in different parts of prearcuate cortex have physiologically distinct signatures. Previous recordings, as well as our own data, suggest that neurons of the prearcuate gyrus reflect formation of decision variables and representation of visual stimuli in the direction discrimination task (Hussar and Pasternak, 2010; Kiani et al., 2014; Kim and Shadlen, 1999; Mante et al., 2013). To determine whether these properties are distributed differentially across prearcuate subnets, we used K-means analysis of the response dissimilarities to divide the recorded units into two mutually exclusive populations, and we projected these populations onto the arrays to visualize their spatial topography (Figures 6A–6C, one example session for each monkey). We then used a logistic model to assess how

well population activity within each subnet predicted trial-to-trial variation in the monkey's upcoming choice and reaction time (RT) (see Supplemental Experimental Procedures). On average, units in subnet-1, the subnet closest to the arcuate sulcus, were more predictive of the monkey's upcoming choice throughout the motion-viewing and delay periods (Figure 6D; *t* test, $p = 0.004$ in the 150 ms window before the Go cue), and thus provided a better representation of the growing decision variable (Shadlen and Newsome, 1996). The subnets were also differentially informative about the monkey's RT (Figure 6E). Although the task was not a RT task, we still observed vari-

ation in the monkey's RT following the Go cue. Interestingly, subnet-1, which was a better predictor of the monkey's choices, was also a better predictor of RT (ANOVA, $p = 0.006$ in the 150 ms window before the Go cue).

The differential representation of the decision-making process by the subnets may appear at odds with our finding above that common noise rather than task-evoked responses underlies the observed topography (Figure 4). We note, however, that the matrix of task-evoked response dissimilarities was weakly but significantly correlated with the matrix of residual response dissimilarities, averaged across all monkeys and sessions (alignment score, ranging from 0.11 ± 0.02 in the postsaccadic epoch to 0.26 ± 0.02 in the motion-viewing epoch; data not shown). That is, pairs of units that show stronger noise correlation also tend to have stronger signal correlation. Thus, knowing the noise correlation of a pair of units offers a weak indication of how the neurons will cooperate in task-related computations (Kenet et al., 2003).

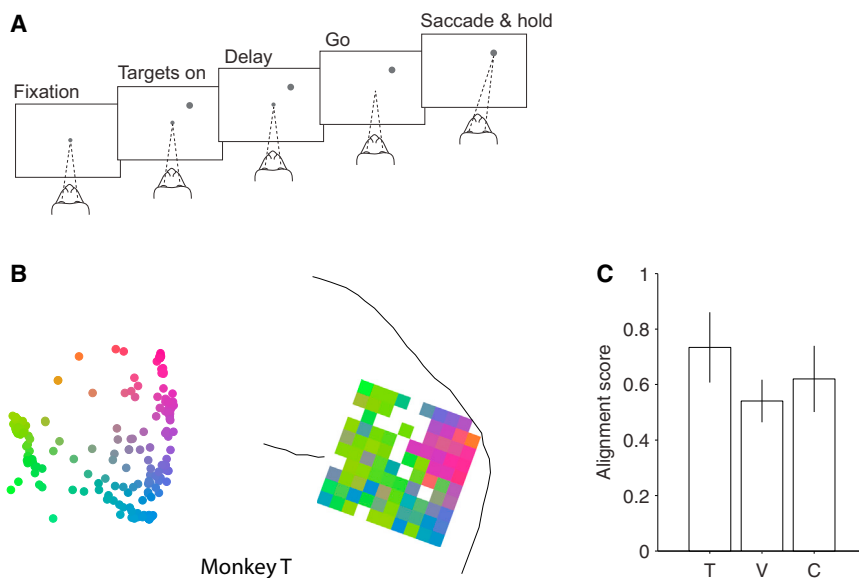


Figure 5. MDS Maps and Spatial Topography Are Invariant to Task Modifications

(A) We recorded neural responses while the monkey performed a second task: visually guided delayed saccade. In this task, after the acquisition of the fixation point by the monkey, a single target was presented on the screen. The monkey made a saccadic eye movement to the target after the Go cue.

(B) The 2D MDS plot and the projected topography on the array for an example session in monkey T. The topography is very similar to that observed in other sessions where the monkey performed a direction discrimination task (e.g., Figure 2F).

(C) The alignment score of the average “electrode-based” dissimilarity matrices (see [Experimental Procedures](#)) across the two tasks. The bars show the alignment score for each monkey. Error bars represent 95% confidence intervals for the alignment between the two dissimilarity matrices.

Task-Induced Variations of Response Dissimilarities

Previous studies have shown that noise correlation can be modulated by spatial attention, context or adaptation (Cohen and Maunsell, 2010; Cohen and Newsome, 2008; Mitchell et al., 2009; Müller et al., 1999), and that response variance is reduced by engagement in a task (Churchland et al., 2010; Purcell et al., 2012). It is important to realize that our findings do not contradict these previous studies. We have shown that *structure* in dissimilarity matrices is largely independent of temporal epoch, including the intertrial interval, but this finding is consistent with modulation in the *overall* level of dissimilarity across trial epochs, as long as the structure is not disturbed. Figure 7A, for example, depicts the complete dissimilarity matrix for each temporal epoch of the experiment illustrated in Figure 2C. The units are segregated by subnet along both the ordinate and the abscissa to facilitate visual comparison of dissimilarity within subnets (upper right and lower left quadrants) and across subnets (upper left and lower right). The overall level of dissimilarity did not change much during the task-related epochs (Constantinidis and Goldman-Rakic, 2002) but varied notably between the intertrial interval (the colors are cooler overall for the intertrial interval) and the five task-related epochs (warm colors). Nevertheless, the structure in the dissimilarity matrices—higher dissimilarity across subnets compared to within subnets—is evident for each epoch (Figure 7B) as well as for the whole-session matrix, and all temporal frequencies (Figure S10). It is this structure that is captured by the MDS plots.

Results from Motor Cortex Are Consistent with Those from Prearcuate Cortex

To extend the scope of our findings, we performed a dissimilarity analysis on data obtained from two multielectrode arrays in a fourth monkey, one placed in the primary motor cortex (M1) and another in the dorsal premotor cortex (PMd) (Figure 8A). The monkey was trained to perform the same direction discrimination task, but reported its choices with reaching movements instead of eye movements. The monkey held its left hand on

the fixation point throughout the trial, and the random dot patch was presented above the fixation point to avoid occlusion by the hand. The task sequence was similar to that illustrated in Figure 1A. We recorded neural activity during task performance as usual (task-engaged period), but we also recorded during extended periods between task blocks (rest periods, 15–60 min). The animal rested calmly in the primate chair in the semi-dark test room during these periods, but without engagement in any behavioral task.

For each recording day, and for all possible pairs of recorded units, we calculated the response dissimilarity matrices separately for the task-engaged periods and the resting periods. The MDS maps from an example experiment (Figure 8B) show clear segregation between the PMd and M1 populations during both periods. For all seven experiments, the response dissimilarity matrices were highly aligned between the task-engaged and rest periods (left-most bar, Figure 8C), confirming the qualitative impression from the example MDS plots. Thus subnet identification is not dependent on attention, arousal states, or specific behavioral events associated with task performance.

Furthermore, three key features of dissimilarity structure in prearcuate cortex were largely preserved in M1 and PMd. First, dissimilarity structure was temporally broadband, persisting across three orders of magnitude of temporal frequency (right bars, Figure 8C). Second, the dissimilarity structure was mainly driven by residual fluctuations of neural activity around the task-evoked mean responses (Figure 8D, left panel) rather than by the means themselves (Figure 8D, right panel). The alignments are generally lower than those for the prearcuate gyrus (Figures 4B and 4C), suggesting a difference across areas. Nonetheless, the difference in the alignment of the residual and task-evoked dissimilarity matrices with the resting period dissimilarity is evident qualitatively and highly significant statistically ($p < 10^{-8}$, three-way ANOVA with response type, session, and epoch as the main factors). Finally, the structure of the dissimilarity matrix was largely preserved across all task epochs,

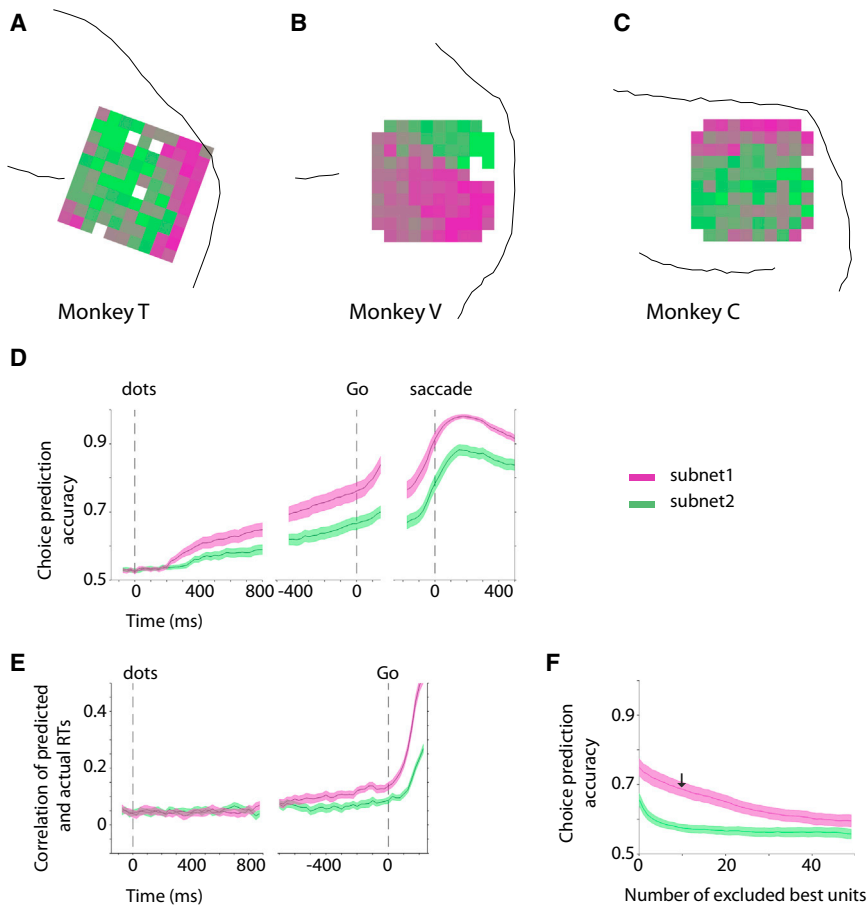


Figure 6. Differential Physiological Properties of the Two Subnets

(A–C) Average layout of the two subnets across the sessions for each monkey. We used K-means clustering to objectively divide the recorded units into two subnets in each session. The subnets were assigned magenta (subnet-1) and green colors (subnet-2) and projected back onto the arrays. The average maps across the sessions are shown for each monkey. The electrodes with in-between colors contributed to different subnets across experiments.

(D) Choice prediction accuracy based on a logistic regression analysis (see [Supplemental Experimental Procedures](#)) of the population responses of subnet-1 and subnet-2.

(E) RT prediction accuracy based on a linear Ridge regression analysis of the population responses of the two subnets. Subnet-1 is a better predictor of both choice and RT. RTs were measured from the Go cue.

(F) Choice predictive responses were more distributed in subnet-1. In each session we ranked individual units of subnet-1 and subnet-2 based on their choice prediction accuracy and then measured the effect of the exclusion of best units on the choice prediction accuracy of the population response. The arrow indicates prediction accuracy of subnet-1 after exclusion of its ten best units. The analysis focuses on the 150 ms window immediately before the Go cue. The shaded areas represent SEM across sessions.

especially around the time of hand movement and during the intertrial interval (data not shown).

Somewhat surprisingly, the resting-state maps were in fact more effective at segregating M1 and PMd, emphasizing our observations that task-related events are not the primary driver of dissimilarity structure and subnet identification (Figure S11). Unfortunately, the number of functional electrodes in these experiments was too low to investigate the spatial topography of potential subnets within a single array, but the parcellation of the cortex into M1 and PMd was clear.

DISCUSSION

We have shown that the recorded neural population in prearcuate gyrus is inhomogeneous and consists of at least two subnets. The responses of neurons within each subnet are more positively correlated with each other and less so with neurons in the other subnet. For each animal, the pattern of correlations across the neural population was largely stable and easily detectable in different tasks and all task epochs, including the intertrial interval. This *pattern* was consistent despite significant variation in the *amplitude* of response correlations across epochs. We also discovered that the functionally defined subnets are spatially segregated in the cortex and are mainly segregated by what is traditionally considered “noise” rather than by the commonly

studied task-evoked responses. These properties hold for motor cortex (M1/PMd) recordings as well.

The basic properties revealed by our subnet analysis—spatial segregation, invariance across behavioral tasks, and adequate definition by spontaneous and task-independent neural activity fluctuations—make our technique an appealing tool for objective parcellation of cortex. It is particularly advantageous in association cortices for two reasons. First, it provides an objective way to group neurons for subsequent analyses; it avoids the “double-dipping” bias caused by emphasizing differences in task-evoked responses following selection of neurons based on the same task-evoked responses (Kriegeskorte et al., 2009). Second, our technique provides easy demarcation of cortical regions of interest in awake, behaving animals. Traditionally, parcellation of cortex has depended heavily on anatomical techniques that cannot be applied in live subjects: cyto-, myelo-, and chemoarchitectonic markers, anterograde and retrograde tracers, and electron microscopy (Amir et al., 1993; Gerbella et al., 2007; Levitt et al., 1993; Rockland and Lund, 1983). More recently, technical advances have enabled cortical parcellation based on fMRI BOLD responses (Power et al., 2011; Vincent et al., 2007), optical imaging, diffusion weighted imaging, and electrocorticography (ECoG) (Hacker et al., 2012; He et al., 2008) in living subjects. To our knowledge, however, our study is the first to do so based on spiking activity in association cortex.

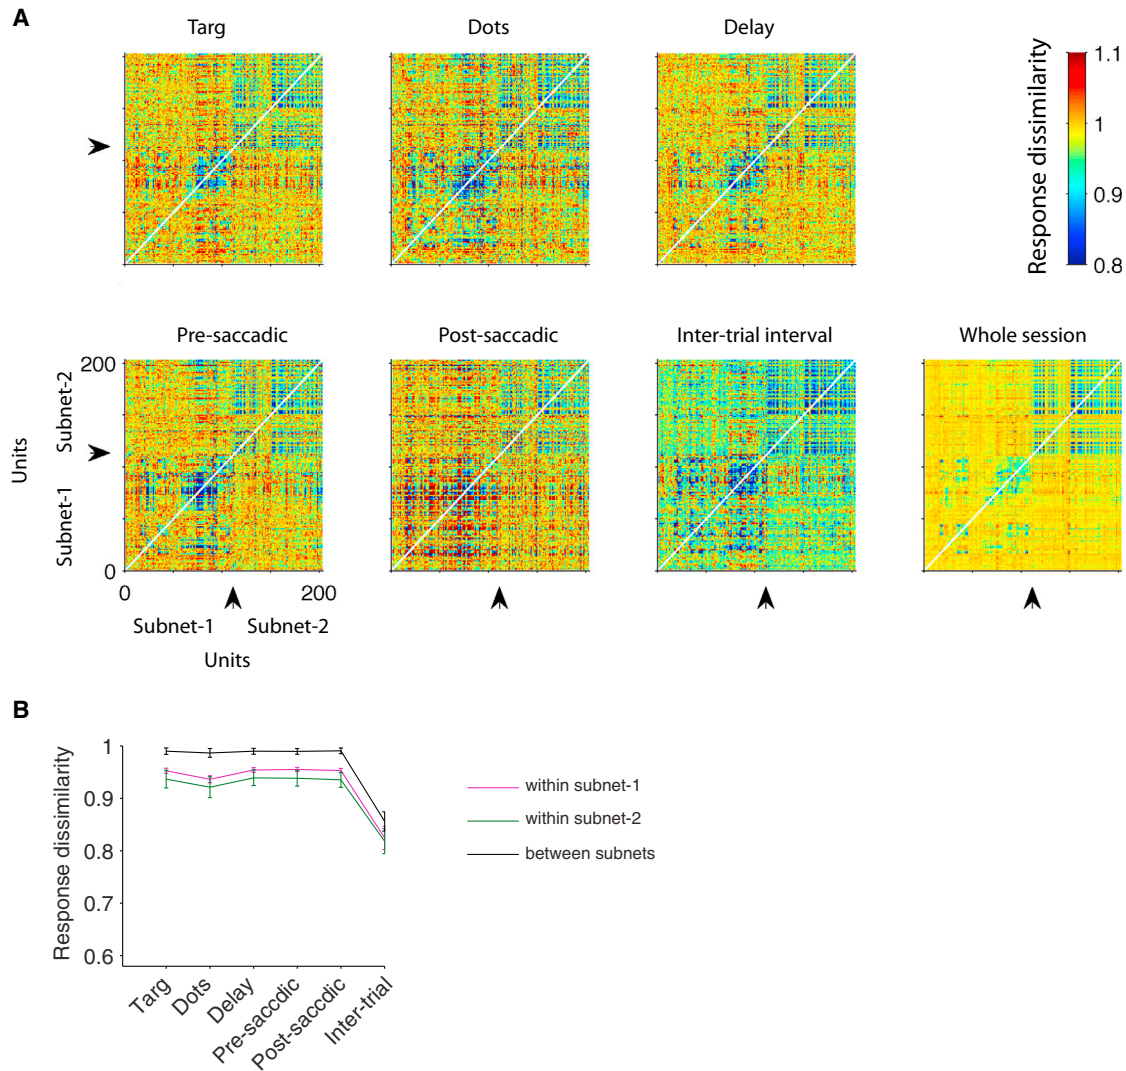


Figure 7. Overall Response Dissimilarity Levels Vary across Task Epochs, but the Structure of the Dissimilarity Matrix Is Stable

(A) The pairwise dissimilarity matrices for all recorded pairs of units in the example session of Figure 2C. Response dissimilarities are measured separately for different task epochs. To facilitate visualization, the units are ordered based on the subnet membership. Arrows indicate the border between the two subnets for this session. The cooler colors during the intertrial interval indicate that dissimilarity is overall lower (response correlation is higher).

(B) Average response dissimilarities within and between the subnets in different task epochs across sessions. Error bars indicate SEM across sessions.

Possible Origins of Correlated Subnet Activity

In the human MRI literature, long-range interareal anatomical connections are emphasized as a source of correlated variability that defines resting state networks (reviewed by Van Dijk et al., 2010), a view that is consistent with the ubiquitous feedforward and feedback pathways connecting cortical areas with each other and with subcortical structures (Felleman and Van Essen, 1991; Markov et al., 2013, 2014). Interareal coordination is particularly striking in the case of cerebocerebellar resting-state circuits for which correlated variability is likely to depend on polysynaptic connections through intermediate structures such as the pons (Habas et al., 2009; Krienen and Buckner, 2009; O'Reilly et al., 2010).

For several reasons, however, we suggest that intrinsic connectivity, especially intra-areal lateral connections, plays a

crucial role in defining the subnets described in this paper. First, stability across tasks and task epochs indicates that the subnets are substantially independent of sensory (or other task-dependent) inputs to the prearcuate gyrus. Although shared input driven by visual stimuli has been shown to modulate the magnitude of pairwise correlations between visual areas (e.g., Jia et al., 2013), our data show that the basic structure of the prearcuate correlation matrix is independent of task epoch and task-evoked responses, and, therefore, unlikely to originate from shared task-related inputs. Similarly, the subnets are independent of motor and decision-related outputs. Among possible task-independent factors, we can rule out slow hemodynamic and neuromodulatory factors as sole causes of correlation structure since the subnets are well

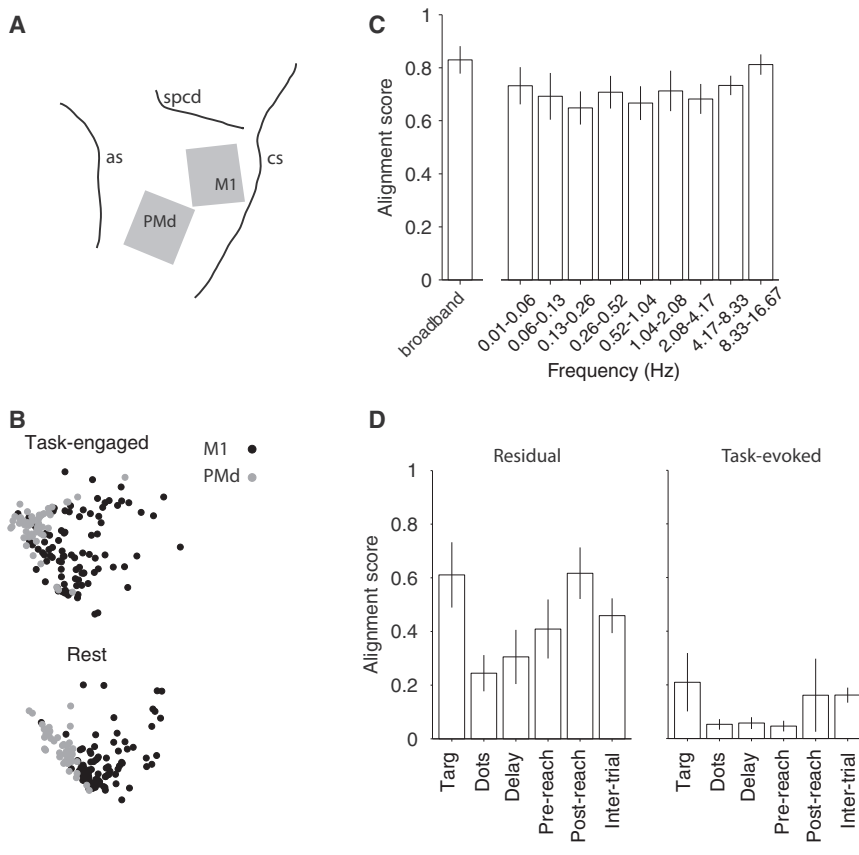


Figure 8. Motor Cortex Data Are Similar to the Prearcuate Data and Extend Our Results to the Resting State

(A) Two multi-electrode arrays were implanted in the left primary motor cortex (M1) and dorsal premotor cortex (PMd) of a monkey trained for a direction discrimination task with reaching movements as the operant response. The gray squares show the array locations with respect to major sulci (as, arcuate sulcus; cs, central sulcus; spcd, superior precentral dimple).

(B) MDS plots of an example session. M1 and PMd are well segregated both during the direction discrimination task and in rest periods between the task-engaged blocks of trials.

(C) Alignment score of the resting and task-engaged dissimilarity matrices. Response dissimilarity matrices were calculated for the combined population across the two arrays. Alignment scores were calculated for the broadband data (unfiltered, leftmost bar) and for the same frequency bands depicted in Figure S6 for the prearcuate data. The matrices are well aligned for the resting and task-engaged periods (left bar), and the signals underlying the alignment are distributed across temporal frequency bands spanning three orders of magnitude (right bars).

(D) Common noise is the main factor underlying the structure of dissimilarity matrices and segregation of M1 and PMd subnets in this analysis. Alignment scores show the correlation between the resting period dissimilarity matrix and the task-evoked (right) and residual dissimilarity matrices (left). Conventions are similar to Figures 4B and 4C.

defined across a wide range of temporal frequencies (Figures S6, S7, and S10).

The second line of evidence is anatomical. A large portion of synapses within a local area of cortex arises from neurons within the same area (intrinsic), not from projections from outside the area (extrinsic). Local horizontal axons and collaterals provide more than half of the excitatory synapses onto pyramidal neurons (Boucsein et al., 2011; Stepanyants et al., 2009), and are thought to coordinate information processing and response dynamics across cortical columns (e.g., Stettler et al., 2002). To the best of our knowledge, lateral connections in 8Ar have not yet been studied, but in the neighboring dorsolateral PFC (areas 46 and 9), lateral connections are organized in patches with dimensions of a few hundred microns to a few millimeters, roughly consistent with the dimensions of the subnets in our study (Levitt et al., 1993).

Finally, modeling studies suggest a prominent role for intrinsic connections in shaping subnets. The topology of connections within a neural circuit molds emergent network dynamics (Buzsáki et al., 2004; Larremore et al., 2011; Ringach, 2009), especially in the absence of external inputs (Galan, 2008). In general, knowledge of network connectivity enables predictions about the correlational structure of the neural responses (Pernice et al., 2012; Trousdale et al., 2012), even though the converse is not true (Kispersky et al., 2011; Sporns, 2012; Trong and Rieke, 2008). Networks that exhibit approximate balance between excitation and inhibition are particularly straightforward in this

respect because response correlations are shaped primarily by the first order connections between neurons rather than by higher-order, polysynaptic chains of intrinsic connections (Trousdale et al., 2012). Networks with balanced excitation and inhibition are likely to be a dominant feature of cortical architecture: they account well for computations known to be carried out in the cortex (Isaacson and Scanziani, 2011), and they produce response variability statistics that correspond closely to those of cortical neurons (Shadlen et al., 1996; van Vreeswijk and Sompolinsky, 1996). Considered together, our current data, coupled with prior anatomical and modeling results, support the role of intrinsic connections as a key determinant of functional subnets defined by dynamic patterns of activity correlation.

Relation of the Subnets to Previous Studies of Area 8Ar

At first glance, the spatial boundary detected by our subnet analysis appears reminiscent of the traditionally defined boundary between the “core” FEF and area 8Ar (Gerbella et al., 2007; Stanton et al., 1989). The FEF lies mostly on the anterior bank of the arcuate sulcus, but can sometimes emerge from the sulcus onto the lip of the prearcuate gyrus (Bruce and Goldberg, 1985; Seidemann et al., 2002), consistent with the close spatial association between subnet 1 and the arcuate sulcus in monkeys T and C (Figures 2F–2H and Figures 6A–6C). This potential association between subnet 1 and the FEF is further suggested by the stronger saccade- and decision-related signals in subnet 1 (Figure 6D).

Arguing against this association, however, are the electrophysiological results from monkey V. For this subject, the functional boundary between the subnets lies directly atop the prearcuate gyrus, and is oriented roughly orthogonally to the nearest point in the arcuate sulcus (Figures 2G and S3). This boundary cannot be reconciled with the standard conceptions of FEF/8Ar. Considering the data as a whole, we suspect that the prearcuate subnets are revealing a functional subdivision separate from the core FEF. With data from only three monkeys, however, this conclusion is tentative and must be considered further in future studies.

Possible Relation of the Subnets to Resting-State fMRI Measurements

The methodology employed in this paper is closely related to the techniques that led to the discovery of resting-state networks in functional imaging studies of the human and monkey brain (Greicius et al., 2003; Vincent et al., 2007), and more recently to large-scale parcellation of the cortex based on “functional connectivity” (e.g., Power et al., 2011). Functional connectivity of two brain voxels is defined as the correlation (or a closely related function) of blood oxygen-level-dependent (BOLD) responses of the voxels. The parcellation techniques, which group together the voxels with covarying BOLD responses, reveal a series of large-scale modules—visual, somatosensory, motor, etc.—consistent with the known large-scale anatomical divisions of the cortex (Honey et al., 2009; Vincent et al., 2007). Interestingly, this large-scale parcellation based on BOLD seems to reflect large-scale electrophysiological properties as well, since the BOLD response fluctuations are closely related to local field potentials and spiking activity within each voxel (Logothetis et al., 2001; Schölvinck et al., 2010). Moreover, the BOLD response correlations across voxels can be mapped to the correlation of the slow cortical potentials in the corresponding locations, as evidenced by ECoG (Hacker et al., 2012; He et al., 2008).

Our results extend these findings in two ways. First, we show that at a small spatial scale, closer to that of cortical columns and intrinsic functional modules, the application of functional connectivity techniques leads to a parcellation similar in robustness and consistency to those observed at much larger spatial scales in functional imaging and ECoG studies. Recall that our recordings were made at 400 micron intervals within a 4 × 4 mm patch of prearcuate gyrus, which roughly corresponds to a single PET or fMRI voxel and is significantly smaller than the spacing of ECoG electrodes. The similarity of results across different techniques hints at shared fundamental principles and a repeated hierarchical organization across different spatial scales (Ganmor et al., 2011).

The second extension relates to the underlying neural events. The fMRI and ECoG signals originate from multiple neural (and possibly nonneural) sources that are difficult to separate from each other (Leopold and Maier, 2012; Logothetis et al., 2001; Moore and Cao, 2008). In contrast, we directly recorded the spiking activity of neurons and avoided the “inverse problem” of decomposing the recorded signal into its constituent events. The underlying neural events in our recordings are thus unitary (spikes) and unambiguous in their location. Our ability to identify a functional boundary from unsupervised analysis of spiking ac-

tivity—even during the intertrial interval (Figures 3 and 8) and during extensive periods of rest from any aspect of task performance (Figures 8 and S11)—suggests a potential neural substrate for the resting state networks identified in functional imaging studies. Networks and parcellation schemes proposed from neuroimaging data will be most compelling if they can be linked definitively to spiking activity of cortical neurons. More definitive links, however, will require simultaneous recordings from a broader expanse of cortex, which can be obtained by implanting multiple microelectrode arrays.

EXPERIMENTAL PROCEDURES

We recorded from populations of neurons in the prearcuate gyrus of three macaque monkeys performing two different tasks: a direction discrimination task and a delayed saccade task. We also recorded from M1 and PMd of a monkey performing a direction discrimination task with reaching responses. All training, surgery, and recording procedures conformed to the National Institutes of Health Guide for the Care and Use of Laboratory Animals and were approved by Stanford University Animal Care and Use Committee.

Behavioral Tasks

Direction Discrimination

Figure 1A illustrates the sequence of events in a single trial of the direction discrimination task for prearcuate recordings. Each trial began with the appearance of a central fixation point (FP; 0.3° diameter) at the center of the monitor. The monkey was required to maintain gaze within $\pm 1.5^\circ$ of FP so long as it was visible on the screen. Eye position was measured with a scleral search coil (CNC Engineering, Seattle). After a short delay, two targets appeared on the monitor. In 21 of 25 sessions the two targets were placed on opposite sides of the screen. In the remaining sessions both targets were placed contralateral to the recorded cortex. After a brief delay the random dots appeared on the screen. The difficulty of the task was controlled by changing the percentage of dots moving coherently in the same direction (motion strength) (Britten et al., 1992; Kiani et al., 2008). The motion strength was chosen randomly on each trial from a set of values that was tailored for each monkey to obtain the full range of performance accuracy from chance (0.5) to nearly perfect (~ 1.0) (Figure 1B). The motion stimulus stayed on the screen for 800 ms and was followed by a variable length delay period (300–1,500 ms, median = 677 ms). The FP disappeared at the end of the delay period (Go cue), signaling the monkey to report the perceived direction of motion with a saccadic eye movement to the corresponding target. The monkey maintained gaze on the target (Figure 1A, “Hold”) until the outcome of the trial was revealed (reward or not, 500–1,000 ms following the operant saccade in most sessions).

Delayed Saccade

After the monkey fixated the FP, a single target appeared on the screen. The location of the target varied from trial to trial and spanned eccentricities up to 25° in each of several directions. The FP disappeared after a variable delay (280–1,300 ms, median = 808 ms), signaling the monkey to make a saccadic eye movement to the target location.

Neural Recording

Multichannel microelectrode arrays (Blackrock Microsystems, Salt Lake City) with 96 electrodes (length = 1.5 mm; spacing = 0.4 mm; impedance ~ 0.5 M Ω) were implanted in the prearcuate gyrus (Figure 1C). The array was positioned between the anterior bank of the concavity of the arcuate sulcus and the posterior tip of principal sulcus in monkeys T and V (Figure 1D). In monkey C the array was placed between the superior branch of arcuate sulcus and dorsal bank of principal sulcus due to anatomical constraints. Neural spike waveforms were saved online (sampling rate, 30 kHz) and sorted offline (Plexon Inc., Dallas). We used customized algorithms to remove recording artifacts that were registered by a large number of electrodes. Also, we merged spike waveform clusters that were judged to be redundant based on waveform shapes, firing rates and interspike intervals. We identified 100–250 single units

and multiunits in each session (median = 219). The spacing of the electrodes was large enough to make recording of the same unit by neighboring electrodes unlikely (Egert et al., 2002). Throughout the paper we use the term “units” to refer to both isolated single neurons and multiunits. All units were retained in our analyses to maximize the spatial coverage of the recorded area and increase the chance of revealing spatial topography.

The direction-discrimination data set included 8, 7, and 11 recording sessions from monkeys T, V, and C, respectively. The delayed-saccade data set included 4, 3, and 2 sessions from the three monkeys. The sessions were chosen based on three factors: large number of trials per session (>1,000), high quality of recordings, and large number of units to provide maximal coverage of the array surface. Relaxing these criteria to increase the number of sessions did not change the results. Although the electrode array remained in a nominally fixed position after surgical insertion, the recorded units frequently changed from one session to another, due, presumably, to small movements of cortex relative to the array.

Behavioral Data Analysis

We fit a cumulative Weibull distribution function to the monkey’s choices,

$$P(\text{cor}) = 1 - 0.5 \times \exp\left(-\left[\frac{C}{\alpha}\right]^\beta\right), \quad \text{Equation 1}$$

where $P(\text{cor})$ is probability correct, C is motion strength, α is psychophysical threshold (the value of C that confers 82% correct responses), and β is a parameter that governs the shape of the function, especially its steepness.

The monkey’s RT was calculated as the delay between the Go cue and saccade initiation. We defined saccade initiation as the time when eye velocity exceeded $15^\circ/\text{s}$.

Neural Data Analysis

For each session, we identified natural physiological groupings of the recorded units based on the dissimilarity of their responses. The response dissimilarity of a pair of units is defined as

$$d_{ij} = 1 - \rho(\vec{r}_i, \vec{r}_j) \quad \text{Equation 2}$$

where \vec{r}_i and \vec{r}_j are the response vectors of units i and j , and $\rho(\vec{r}_i, \vec{r}_j)$ is Pearson’s correlation. The response dissimilarity, therefore, reflects covariation of neural responses and can take any value between 0 (perfect correlation) and 2 (perfect anticorrelation). For the whole-session analyses (e.g., Figure 2) we defined the neural response vector for each unit in 30 ms nonoverlapping bins from the beginning of the session to its end, independent of task epochs, visual stimuli, and the monkey’s behavior. The neural response vector varied in other analyses, as explained below. Using Equation 2, we calculated dissimilarity for all possible pairs of units in a given experiment; throughout the paper, we refer to this set of metrics as the *dissimilarity matrix* for the corresponding experiment.

To visualize the relationship between units and investigate their grouping, we applied MDS to the dissimilarity matrix. MDS creates a low-dimensional representation that retains the pairwise relationships as much as possible. Each point on our MDS maps (Figure 2A) represents a recorded unit. All units were included in the analysis to maximize the coverage of the recording array. The Euclidean distance between units on the MDS map reflects how the neural responses of those units covary—shorter distances suggest higher correlations. We used a nonlinear MDS technique (Isomap), (De’ath, 1999; Tenenbaum and Freeman, 2000), but we obtained similar results with other MDS methods. Our 2D MDS maps captured 49.2%–82.9% of variance of the dissimilarity matrix (mean = 61.2%). Figure 2B illustrates the average Scree plot across sessions.

To explore the spatial relationships of units on the cortical surface we chose a unique color for each unit based on its location in the 2D MDS map and a spatially smooth 2D color map (Figures 2C–2E). This color was then assigned to the location of the electrode that recorded the unit (Figures 2F–2H and S3). The locations with similar colors, therefore, recorded units that were close to each other on the MDS map. When more than one unit was recorded on a single electrode, the colors for the individual units were averaged, and the color corresponding to the average was assigned to that electrode location.

For the epoch-based analysis (Figures 3 and S4), we only used the neural responses coming from a single trial epoch. Six different epochs are used in this paper: target onset (50–300 ms after target onset), motion viewing (300–800 ms after motion onset), delay period (50–250 ms after motion offset), presaccadic period (250–50 ms before the saccade), postsaccadic period (50–250 ms after the saccade), and intertrial interval (150–400 ms after the eye left the target window). The response intervals were chosen to be representative of the response dynamics of the recorded units; our results do not depend strongly on the exact temporal boundaries of these intervals. Within each epoch the responses could be measured as the total spike count or as a vector of spike counts in successive 30 ms windows that tile the epoch window. The results do not critically depend on which option was used. To calculate epoch-specific dissimilarity metrics, data from a particular epoch were concatenated together across all trials of a session, omitting data from all other epochs. Dissimilarity was then calculated on the concatenated data using Equation 2.

To visualize the match of the spatial topographies across epochs, we created MDS maps independently for each epoch (Figure 3), but borrowed the color of the units from the whole-session MDS map for that monkey (Figure 2C). Thus, clustering of units with similar colors in the new maps indicates a good match of an epoch map to the original whole-session map. MDS maps were created only for visualization of the data. To quantify the alignment of dissimilarity matrices we calculated their correlation (alignment score) and used Mantel’s test (Mantel, 1967) to assess the significance of the correlation. Exclusion of the aborted and/or error trials did not significantly influence the conclusions.

The high correlation of dissimilarity matrices across epochs (Results) suggests minimal influence of task parameters on response dissimilarities. We looked for the source of these effects by breaking the responses of individual units into two components: a task-evoked component calculated as the average response across all trials with similar motion direction, motion strength, and choice; and a residual component calculated as the fluctuation of the response around that mean on each trial. Similar results were obtained if the residual responses were normalized by the standard deviation of responses of the trials with similar choice, motion direction, and strength (Bair et al., 2001). Task-evoked and residual responses were calculated for the epoch durations explained above. To ensure reliability, we excluded conditions with fewer than 30 trials. The excluded conditions consisted mainly of erroneous choices on medium- and high-coherence trials. We recalculated the epoch-based dissimilarity matrices for each response component (task-evoked and residual) and measured their alignments to the whole-session dissimilarity matrix (Figure 4).

To test whether the dissimilarity matrices (and thus potential physiological groupings) were consistent across tasks, we measured the alignment between response dissimilarity matrices in the direction discrimination and delayed saccade tasks (Figure 5). The data sets for the two tasks were collected in different recording sessions in order to maximize the trial counts per data set. Because the recorded units could change from one session to the next, our comparison of dissimilarity across the two tasks was limited in accuracy. To make the comparison as accurate as possible, we reduced between-session variation by first calculating an average dissimilarity matrix across all sessions of a particular task before measuring the alignment of dissimilarity matrices between tasks. Specifically, we first calculated the average response dissimilarity of all pairs of units recorded by each pair of electrodes in a given session (electrode-based dissimilarity). For each monkey, each task, and each pair of electrodes, we then averaged these dissimilarities across all sessions to provide the best estimate for the dissimilarity of the neural population recorded by each pair of electrodes. Finally, we measured the alignment of these average response dissimilarities between the two tasks.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures and 11 figures and can be found with this article at <http://dx.doi.org/10.1016/j.neuron.2015.02.014>.

ACKNOWLEDGMENTS

We thank Daniel Kimmel, Valerio Mante, Vince McGinty, Jonathan Winawer, Brian Wandell, Elad Schneidman, and Kanaka Rajan for useful discussions. Also, we thank Jessica Powell, Jamie Sanders, Sania Fong, and Julian Brown for animal care and technical support. This research was funded by Howard Hughes Medical Institute, the Air Force Research Laboratory (agreement number FA9550-07-1-0537), a Berry Post-doctoral Fellowship to R.K., and Simons Collaboration on the Global Brain.

Received: March 30, 2014

Revised: October 6, 2014

Accepted: January 31, 2015

Published: February 26, 2015

REFERENCES

- Abe, H., and Lee, D. (2011). Distributed coding of actual and hypothetical outcomes in the orbital and dorsolateral prefrontal cortex. *Neuron* 70, 731–741.
- 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.
- Badre, D., and D'Esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical? *Nat. Rev. Neurosci.* 10, 659–669.
- 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.
- Boucsein, C., Nawrot, M.P., Schnepel, P., and Aertsen, A. (2011). Beyond the cortical column: abundance and physiology of horizontal connections imply a strong role for inputs from the surround. *Front. Neurosci.* 5, 32.
- Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* 12, 4745–4765.
- Bruce, C.J., and Goldberg, M.E. (1985). Primate frontal eye fields. I. Single neurons discharging before saccades. *J. Neurophysiol.* 53, 603–635.
- Bruce, C.J., Goldberg, M.E., Bushnell, M.C., and Stanton, G.B. (1985). Primate frontal eye fields. II. Physiological and anatomical correlates of electrically evoked eye movements. *J. Neurophysiol.* 54, 714–734.
- Buzsáki, G., Geisler, C., Henze, D.A., and Wang, X.J. (2004). Interneuron Diversity Series: circuit complexity and axon wiring economy of cortical interneurons. *Trends Neurosci.* 27, 186–193.
- Carmichael, S.T., and Price, J.L. (1994). Architectonic subdivision of the orbital and medial prefrontal cortex in the macaque monkey. *J. Comp. Neurol.* 346, 366–402.
- Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13, 369–378.
- Cohen, M.R., and Maunsell, J.H. (2010). A neuronal population measure of attention predicts behavioral performance on individual trials. *J. Neurosci.* 30, 15241–15253.
- Cohen, M.R., and Newsome, W.T. (2008). Context-dependent changes in functional circuitry in visual area MT. *Neuron* 60, 162–173.
- 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.
- De'ath, G. (1999). Extended dissimilarity: a method of robust estimation of ecological distances from high beta diversity data. *Plant Ecol.* 144, 191–199.
- Egert, U., Heck, D., and Aertsen, A. (2002). Two-dimensional monitoring of spiking networks in acute brain slices. *Exp. Brain Res.* 142, 268–274.
- Engel, S.A., Rumelhart, D.E., Wandell, B.A., Lee, A.T., Glover, G.H., Chichilnisky, E.J., and Shadlen, M.N. (1994). fMRI of human visual cortex. *Nature* 369, 525.
- Felleman, D.J., and Van Essen, D.C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex* 1, 1–47.
- Galan, R.F. (2008). On how network architecture determines the dominant patterns of spontaneous neural activity. *PLoS ONE*. <http://dx.doi.org/10.1371/journal.pone.0002148>.
- Ganmor, E., Segev, R., and Schneidman, E. (2011). Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *Proc. Natl. Acad. Sci. USA* 108, 9679–9684.
- Gerbella, M., Belmalih, A., Borra, E., Rozzi, S., and Luppino, G. (2007). Multimodal architectonic subdivision of the caudal ventrolateral prefrontal cortex of the macaque monkey. *Brain Struct. Funct.* 212, 269–301.
- Greicius, M.D., Krasnow, B., Reiss, A.L., and Menon, V. (2003). Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proc. Natl. Acad. Sci. USA* 100, 253–258.
- Habas, C., Kamdar, N., Nguyen, D., Prater, K., Beckmann, C.F., Menon, V., and Greicius, M.D. (2009). Distinct cerebellar contributions to intrinsic connectivity networks. *J. Neurosci.* 29, 8586–8594.
- Hacker, C.D., Snyder, A.Z., Sharma, M., Bundy, D.T., Daitch, A.L., Szrama, N., Pahwa, M., Gaona, C.M., Corbetta, M., and Leuthardt, E.C. (2012). The electrophysiology of resting-state fMRI networks. In *Society for Neuroscience Meeting*, New Orleans.
- He, B.J., Snyder, A.Z., Zempel, J.M., Smyth, M.D., and Raichle, M.E. (2008). Electrophysiological correlates of the brain's intrinsic large-scale functional architecture. *Proc. Natl. Acad. Sci. USA* 105, 16039–16044.
- Honey, C.J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.P., Meuli, R., and Hagmann, P. (2009). Predicting human resting-state functional connectivity from structural connectivity. *Proc. Natl. Acad. Sci. USA* 106, 2035–2040.
- Hubel, D.H., and Livingstone, M.S. (1987). Segregation of form, color, and stereopsis in primate area 18. *J. Neurosci.* 7, 3378–3415.
- Hussar, C.R., and Pasternak, T. (2009). Flexibility of sensory representations in prefrontal cortex depends on cell type. *Neuron* 64, 730–743.
- Hussar, C., and Pasternak, T. (2010). Trial-to-trial variability of the prefrontal neurons reveals the nature of their engagement in a motion discrimination task. *Proc. Natl. Acad. Sci. USA* 107, 21842–21847.
- Isaacson, J.S., and Scanziani, M. (2011). How inhibition shapes cortical activity. *Neuron* 72, 231–243.
- Jia, X., Tanabe, S., and Kohn, A. (2013). γ and the coordination of spiking activity in early visual cortex. *Neuron* 77, 762–774.
- Katz, L.C., Gilbert, C.D., and Wiesel, T.N. (1989). Local circuits and ocular dominance columns in monkey striate cortex. *J. Neurosci.* 9, 1389–1399.
- Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., and Arieli, A. (2003). Spontaneously emerging cortical representations of visual attributes. *Nature* 425, 954–956.
- Kiani, R., Hanks, T.D., and Shadlen, M.N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci.* 28, 3017–3029.
- Kiani, R., Cueva, C.J., Reppas, J.B., and Newsome, W.T. (2014). Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials. *Curr. Biol.* 24, 1542–1547.
- Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* 2, 176–185.
- Kispersky, T., Gutierrez, G.J., and Marder, E. (2011). Functional connectivity in a rhythmic inhibitory circuit using Granger causality. *Neural Sys. Circ.* 1, 9.
- Kohn, A., Zandvakili, A., and Smith, M.A. (2009). Correlations and brain states: from electrophysiology to functional imaging. *Curr. Opin. Neurobiol.* 19, 434–438.
- Kriegeskorte, N., Simmons, W.K., Bellgowan, P.S., and Baker, C.I. (2009). Circular analysis in systems neuroscience: the dangers of double dipping. *Nat. Neurosci.* 12, 535–540.

- Krienen, F.M., and Buckner, R.L. (2009). Segregated fronto-cerebellar circuits revealed by intrinsic functional connectivity. *Cereb. Cortex* *19*, 2485–2497.
- Larremore, D.B., Shew, W.L., Ott, E., and Restrepo, J.G. (2011). Effects of network topology, transmission delays, and refractoriness on the response of coupled excitable systems to a stochastic stimulus. *Chaos* *21*, 025117.
- 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.
- Lennert, T., and Martinez-Trujillo, J.C. (2013). Prefrontal neurons of opposite spatial preference display distinct target selection dynamics. *J. Neurosci.* *33*, 9520–9529.
- Leopold, D.A., and Maier, A. (2012). Ongoing physiological processes in the cerebral cortex. *Neuroimage* *62*, 2190–2200.
- LeVay, S., Hubel, D.H., and Wiesel, T.N. (1975). The pattern of ocular dominance columns in macaque visual cortex revealed by a reduced silver stain. *J. Comp. Neurol.* *159*, 559–576.
- 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.K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature* *412*, 150–157.
- 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.
- Mantel, N. (1967). Ranking procedures for arbitrarily restricted observation. *Biometrics* *23*, 65–78.
- Markov, N.T., Ercsey-Ravasz, M., Van Essen, D.C., Knoblauch, K., Toroczkai, Z., and Kennedy, H. (2013). Cortical high-density counterstream architectures. *Science* *342*, 1238406.
- Markov, N.T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., Lamy, C., Misery, P., Giroud, P., Ullman, S., et al. (2014). Anatomy of hierarchy: feedforward and feedback pathways in macaque visual cortex. *J. Comp. Neurol.* *522*, 225–259.
- Mishkin, M., Ungerleider, L.G., and Macko, K.A. (1983). Object vision and spatial vision—2 cortical pathways. *Trends Neurosci.* *6*, 414–417.
- Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2009). Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* *63*, 879–888.
- Monosov, I.E., and Hikosaka, O. (2012). Regionally distinct processing of rewards and punishments by the primate ventromedial prefrontal cortex. *J. Neurosci.* *32*, 10318–10330.
- Moore, C.I., and Cao, R. (2008). The hemo-neural hypothesis: on the role of blood flow in information processing. *J. Neurophysiol.* *99*, 2035–2047.
- Müller, J.R., Metha, A.B., Krauskopf, J., and Lennie, P. (1999). Rapid adaptation in visual cortex to the structure of images. *Science* *285*, 1405–1408.
- O'Reilly, J.X., Beckmann, C.F., Tomassini, V., Ramnani, N., and Johansen-Berg, H. (2010). Distinct and overlapping functional zones in the cerebellum defined by resting state functional connectivity. *Cereb. Cortex* *20*, 953–965.
- O'Scalaidhe, S.P., Wilson, F.A., and Goldman-Rakic, P.S. (1997). Areal segregation of face-processing neurons in prefrontal cortex. *Science* *278*, 1135–1138.
- Pernice, V., Staude, B., Cardanobile, S., and Rotter, S. (2012). Recurrent interactions in spiking networks with arbitrary topology. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* *85*, 031916.
- Petrides, M., and Pandya, D.N. (1999). Dorsolateral prefrontal cortex: comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns. *Eur. J. Neurosci.* *11*, 1011–1036.
- Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A., Vogel, A.C., Laumann, T.O., Miezin, F.M., Schlaggar, B.L., and Petersen, S.E. (2011). Functional network organization of the human brain. *Neuron* *72*, 665–678.
- Preuss, T.M. (2007). Evolutionary specializations of primate brain systems. In *Primate Origins: Adaptations and Evolution* M, J. Ravosa and M. Dagasto, eds. (New York: Springer), pp. 625–675.
- Purcell, B.A., Heitz, R.P., Cohen, J.Y., and Schall, J.D. (2012). Response variability of frontal eye field neurons modulates with sensory input and saccade preparation but not visual search salience. *J. Neurophysiol.* *108*, 2737–2750.
- Ringach, D.L. (2009). Spontaneous and driven cortical activity: implications for computation. *Curr. Opin. Neurobiol.* *19*, 439–444.
- Robinson, D.A., and Fuchs, A.F. (1969). Eye movements evoked by stimulation of frontal eye fields. *J. Neurophysiol.* *32*, 637–648.
- Rockland, K.S., and Lund, J.S. (1983). Intrinsic laminar lattice connections in primate visual cortex. *J. Comp. Neurol.* *216*, 303–318.
- Romanski, L.M., and Goldman-Rakic, P.S. (2002). An auditory domain in primate prefrontal cortex. *Nat. Neurosci.* *5*, 15–16.
- Schall, J.D. (1997). Visuomotor areas of the frontal lobe. In *Cerebral Cortex*, K.S. Rockland, A.J. Peters, and J. Kaas, eds. (New York: Plenum Press).
- Schölvinck, M.L., Maier, A., Ye, F.Q., Duyn, J.H., and Leopold, D.A. (2010). Neural basis of global resting-state fMRI activity. *Proc. Natl. Acad. Sci. USA* *107*, 10238–10243.
- Seidemann, E., Arieli, A., Grinvald, A., and Slovin, H. (2002). Dynamics of depolarization and hyperpolarization in the frontal cortex and saccade goal. *Science* *295*, 862–865.
- Shadlen, M.N., and Newsome, W.T. (1996). Motion perception: seeing and deciding. *Proc. Natl. Acad. Sci. USA* *93*, 628–633.
- Shadlen, M.N., Britten, K.H., Newsome, W.T., and Movshon, J.A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* *16*, 1486–1510.
- Smith, M.A., and Kohn, A. (2008). Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* *28*, 12591–12603.
- Sporns, O. (2012). From simple graphs to the connectome: networks in neuroimaging. *Neuroimage* *62*, 881–886.
- Stanton, G.B., Deng, S.Y., Goldberg, M.E., and McMullen, N.T. (1989). Cytoarchitectural characteristic of the frontal eye fields in macaque monkeys. *J. Comp. Neurol.* *282*, 415–427.
- Stepanyants, A., Martinez, L.M., Ferecskó, A.S., and Kisvárdy, Z.F. (2009). The fractions of short- and long-range connections in the visual cortex. *Proc. Natl. Acad. Sci. USA* *106*, 3555–3560.
- Stettler, D.D., Das, A., Bennett, J., and Gilbert, C.D. (2002). Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neuron* *36*, 739–750.
- Tenenbaum, J.B., and Freeman, W.T. (2000). Separating style and content with bilinear models. *Neural Comput.* *12*, 1247–1283.
- Trong, P.K., and Rieke, F. (2008). Origin of correlated activity between parasol retinal ganglion cells. *Nat. Neurosci.* *11*, 1343–1351.
- Trousdale, J., Hu, Y., Shea-Brown, E., and Josić, K. (2012). Impact of network structure and cellular response on spike time correlations. *PLoS Comput. Biol.* *8*, e1002408.
- Tsao, D.Y., Schweers, N., Moeller, S., and Freiwald, W.A. (2008). Patches of face-selective cortex in the macaque frontal lobe. *Nat. Neurosci.* *11*, 877–879.
- Tsodyks, M., Kenet, T., Grinvald, A., and Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* *286*, 1943–1946.
- Van Dijk, K.R., Hedden, T., Venkataraman, A., Evans, K.C., Lazar, S.W., and Buckner, R.L. (2010). Intrinsic functional connectivity as a tool for human connectomics: theory, properties, and optimization. *J. Neurophysiol.* *103*, 297–321.

- Van Essen, D.C., Newsome, W.T., and Maunsell, J.H. (1984). The visual field representation in striate cortex of the macaque monkey: asymmetries, anisotropies, and individual variability. *Vision Res.* 24, 429–448.
- van Vreeswijk, C., and Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274, 1724–1726.
- Vincent, J.L., Patel, G.H., Fox, M.D., Snyder, A.Z., Baker, J.T., Van Essen, D.C., Zempel, J.M., Snyder, L.H., Corbetta, M., and Raichle, M.E. (2007). Intrinsic functional architecture in the anaesthetized monkey brain. *Nature* 447, 83–86.
- Wiesel, T.N., and Hubel, D.H. (1974). Ordered arrangement of orientation columns in monkeys lacking visual experience. *J. Comp. Neurol.* 158, 307–318.
- Zeki, S., and Shipp, S. (1988). The functional logic of cortical connections. *Nature* 335, 311–317.