Supplemental Experimental Procedures

Length constant of dissimilarity

It has been reported previously that the correlation of neural responses depends on spatial distance between the recorded units (Constantinidis and Goldman-Rakic, 2002; Leavitt et al., 2013; Smith and Kohn, 2008). We calculated the length constant of dissimilarity by fitting an exponential function to the curves that relate dissimilarity between units to inter-electrode distance (Fig. S2):

$$d_{ij} = d_{\infty} - \beta \exp\left(-\frac{L_{ij}}{\tau}\right) \qquad \text{Eq. S1}$$

where β, d_{∞}, and τ are model parameters, and L_{ij} is the physical distance between recording electrodes on the array. d_{∞} defines the asymptotic dissimilarity for widely separated units. τ is the length constant of the equation and defines the distance at which d_{ij} increases by 63% of the difference between d_0 and d_{∞}, where d_0 is the expected dissimilarity of units recorded on the same electrode.

Temporal structure of the signal that underlies the subnets

We explored the dependence of the response dissimilarity matrix on temporal frequency (Fig. S6 and S7) by creating 9 different frequency bands: 0.01-0.06 Hz, 0.06-0.13 Hz, 0.13-0.26 Hz, 0.26-0.52 Hz, 0.52-1.04 Hz, 1.04-2.08 Hz, 2.08-4.17 Hz, 4.17-8.33 Hz, 8.33-16.67 Hz. Note that 16.67 Hz is the Nyquist frequency for our 30 ms response window. We calculated the dissimilarity matrix for each frequency band using two different methods: by applying a 3rd order Butterworth bandpass filter to the response vector of individual units before calculating correlations; and by performing Fourier transform on the response vector and zeroing coefficients outside the target band. The end results of the two methods were highly compatible. We present the former in this paper because it does not generate 'ripples' in the firing rate vectors. For each frequency band, we calculated the MDS maps for visual assessment, and we measured quantitatively the alignment of the filtered and unfiltered dissimilarity matrices.

Physiological properties of the subnets

Our initial visual inspections of the MDS maps indicated the presence of inhomogeneity in the neural population and thus the possibility that the recorded population could be divided into distinct sub-networks ('subnets') of physiologically related units. We used K-means clustering (MacQueen, 1967) to objectively divide the population into two subnets. Further, we tested the significance of this division using SigClust (Liu et al., 2008). The significance of the divisions was tested in 2-dimensional MDS projections (see Results). Our use of K-means to specify two significantly different populations was conservative. In some experiments, there are indications that the recorded neural population may in fact comprise more than two spatially segregated clusters. Because of

the high dimensionality of the dataset, however, proving the existence of more than two clusters is statistically implausible. Demonstrating a finer grain clustering in area 8Ar will require higher density recording techniques.

We compared the physiological response properties of the subnets using three different metrics (Fig. 6): the prediction accuracy for the monkey's choices based on the population responses in each subnet, the prediction accuracy for the monkey's reaction time, and the frequency of choice predictive and choice 'postdictive' units in each subnet.

We employed a logistic regression model to predict the monkey's choice based on the neural responses:

$$\text{logit}\left[P(choice)\right] = \beta_0 + \sum_{i=1}^{n} \beta_i r_i \qquad \text{Eq. S2}$$

where r_i is the response of unit *i*, *n* is the number of units in the subnet, and the β coefficients are model parameters. The model was cross-validated by fitting to 90% of trials and predicting the remaining 10% in each session (10-fold cross-validation). To generate the time-varying accuracy functions in Figure 6, we first fit the model separately to the neural responses measured in a 150 ms window centered at each moment in time. Then, we averaged the temporal profiles of prediction accuracy across sessions. The p-value for the difference of choice prediction accuracy of the two subnets was obtained by running a t-test in the time window immediately before the Go cue.

We employed a linear Ridge regression model to predict the monkey's RT:

$$RT = \beta_0 + \sum_{i=1}^{n} \beta_i r_i \qquad \text{Eq. S3}$$

where, once again, r_i is the response of unit *i*, *n* is the number of units in the subnet, and the β coefficients are fitted model parameters (Ridge parameter=1000). Again, we employed a 10-fold cross-validation method. The predictions were performed separately for choices to each target location, and the goodness of the predictions was measured by calculating the correlation of predicted and actual RTs. A 150 ms sliding window was used for the time varying plots in Figure 6. To test whether one subnet predicted RT significantly better than the other, we performed a two-way, nested ANOVA with subnet and target location as main factors and the correlation of predicted and actual RTs as the dependent variable. The test was performed on responses measured during the time window immediately prior to the Go cue.

Finally, we compared 'postdictive' encoding of target choice by the two subnets. 'Postdictive' encoding refers to signals that arise during the temporal interval *after* the operant choice is made but *before* the reward is delivered. These signals, which appear to comprise a memory trace of the recently made eye movement, have been described previously by other groups (Tsujimoto et al., 2010) and in these prearcuate data by Reppas and colleagues (Reppas and Newsome, 2008). Postdictive responses are frequently the strongest signals carried by the prearcuate neurons. We assessed postdictive selectivity during a 250 ms window after the operant saccade and immediately before reward delivery (t-test, $p<0.05$). A χ^2-test was used to test the

difference in the distribution of choice-related units in the two subnets. The statistical tests in this paper assume independence across sessions. The sessions from different subjects are indeed independent of each other but the sessions from the same subject overlap because some of the units are likely to be the same from session to session. This overlap violates the independence assumption and calls for development of new techniques for quantifying overlap and upward correction of p-values. However, due to the strength of our results and their consistency across the three animals, such a correction is highly unlikely to influence our conclusions.

Some care must be taken to distinguish predictive and postdictive responses. A predictive cell whose activity is informative about an upcoming saccade (t-test, $p<0.05$ for firing rates in a 250 ms window immediately before the operant saccade) can *appear* to encode saccade direction postdictively if it also generates predictive activity before the 'exit' saccade that follows reward delivery at the end of the trial. Exit saccades at the end of the trial typically return the animal's gaze to the location of the fixation point—a saccade exactly opposite in direction and amplitude to the operant saccade. Thus the signal during the postdictive interval will appear to change polarity, or 'flip', in comparison to the selectivity immediately prior to the operant saccade. Following the conservative procedure of Reppas and colleagues (Reppas and Newsome, 2008), we excluded all 'flipper' cells from our population analysis of postdictive signals. For units exhibiting postdictive signals, we quantified the size of the effect by measuring the area under the ROC curve (Green and Swets, 1966) computed from responses following saccades to each of the two targets.

Supplemental Results

The signal that underlies subnets is temporally broadband

Results presented in the main text show that the subnets do not arise primarily from task-related events or processes, nor from task engagement *per se* (the subnets exist during the inter-trial interval as well). We can also rule out trivial factors such as spiking noise, which results primarily from failures of transmission at individual synapses and is therefore independent across neurons (Rieke et al., 1997). The driving factor underlying the subnets must be some fundamental aspect of neuronal circuitry within the prearcuate gyrus, intrinsic connectivity being a prime candidate (Barnett et al., 2009; Galan, 2008; Pernice et al., 2012; Robinson, 2012; Sporns, 2011; Sporns et al., 2000). We cannot, however, rule out *a priori* a role for slower processes such as hemodynamic (Moore and Cao, 2008) or neuromodulatory differences between subnets. To shed light on these possibilities, we explored how response dissimilarities vary across temporal frequency bands.

We recalculated the dissimilarities within nine temporal frequency bands, from 0.01 Hz to 16.7 Hz (see Supplemental Experimental Procedures; 16.7 Hz is the Nyquist frequency for our 30 ms response window), and measured their alignment to whole-session response dissimilarities. The MDS plots for different temporal frequency bands reveal clustering that is similar to whole-session clustering (Fig. S6A), an impression that is

confirmed quantitatively by the alignment scores (Fig. 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. Note that the *magnitude* of correlation of pairs of neurons was dependent on temporal frequency (Fig. S10), as expected from past research (e.g. He et al., 2010b; Smith and Kohn, 2008), even though the *structure* of the correlation matrix across the population was largely independent of temporal frequency (Fig. S6, S7, and S10). The magnitude of the pairwise correlations can vary between frequency bands while the overall pattern of pairwise correlations is maintained. That is, if pair A is more strongly correlated than pair B, that order is maintained across frequency bands even if the absolute magnitude of the correlations change. Preservation of high alignment scores and color clusters in the MDS plots means that, within each frequency band, spatial topography on the array resembles the whole-session topography illustrated in Figure 2F (Fig. S7).

The peak alignment in the delta band was not due to an overabundance of signal power in that band. The power spectrum of response magnitudes (in 30 ms bins across an entire experiment) had a $1/f^{\beta}$ shape with greatest power at the lowest frequencies, as expected from prior results (Bedard et al., 2006; He et al., 2010a; Leopold and Logothetis, 2003; Teich et al., 1997).

Additional differences in physiological properties of subnets

Not only were the responses of subnet-1 better predictors of choice and RT (see main text), subnet-1 also exhibited a more distributed neural representation of choice information. For each subnet, we tested whether choice predictive power was spread across many units, or was dependent on a select few predictive units. To test this, we excluded each unit from the population one-by-one in descending order of predictive power, and used the logistic model to predict the monkey's choice from the rest of the population. For this analysis we focused on the 150 ms window immediately before the Go cue. Exclusion of a small number of units from subnet-2 caused choice predictive accuracy to drop rapidly, whereas subnet-1 was significantly more resilient to the exclusion of its best units (Fig. 6F). Moreover, subnet-1 yielded better predictive activity with its 10 best units excluded than did subnet-2 with no units excluded (arrow, Fig. 6F).

In addition to assessing choice predictive activity, we also measured choice 'postdictive' activity— a retrospective coding of choice during an interval that follows the operant saccade but precedes the time of potential reward delivery (Tsujimoto et al., 2009, 2010) (see Supplemental Experimental Procedures). Choice-related neurons were distributed differentially across the subnets. More than half of the units of subnet-1 (51%) exhibited choice-related responses (either predictive or postdictive), whereas only 26% of the units in subnet-2 yielded choice-related responses (χ^2-test, $p<10^{-8}$). However, a larger fraction of choice-related units in subnet-2 showed postdictive selectivity as compared to subnet-1 (78% vs. 68%; χ^2-test, $p=8\times10^{-6}$).

Additional control analyses of neural data

We performed additional control analyses to test whether the structure of the dissimilarity matrix for a given experiment could be explained simply by differences in overall firing rates across units or the recorded cortical layers within an experiment. The first concern is that if firing rates vary considerably across the array, the known relationship between firing rates and the magnitude of correlated variability between neurons (de la Rocha et al., 2007) might provide a trivial explanation for the structure of dissimilarity matrix. This was not the case. For the example experiment in Figs. 2C-F, 3 and 4 we calculated a matrix of the difference in average firing rate between all pairs of units on the array, and used our alignment score to compare the matrix of firing rate differences to the matrix of correlation dissimilarities. For all trial epochs, the alignment score was very small (< 0.06, data not shown), demonstrating that variation in firing rate, *per se*, cannot explain the structure in the dissimilarity matrix.

The data from monkey V allow us to reject an alternative interpretation of the subnet analysis, related to laminar organization. The arrays were implanted on relatively flat cortex of the prearcuate gyrus, and the 1.5 mm long electrodes would generally be expected to sample cells from lower layer 3, layer 4, and upper layer 5. Along the edge of the array nearest to the arcuate sulcus, however, some electrodes may have preferentially sampled layer 2 and upper layer 3 due to the changing orientation of the cortex as it folds into the arcuate sulcus. If the upper layers have different noise characteristics than the lower layers, the pattern of subnets observed in monkeys T and C might result from this differential sampling. However, the subnet boundary in monkey V, which lies squarely atop the prearcuate gyrus, argues strongly against this interpretation.

This conclusion is further supported by similarity of indices that have been suggested to depend on laminar location of recorded neurons: overall firing rates (see above) and magnitude of pairwise correlations. If different parts of the array are recording from different layers of cortex one might expect to see a topographical organization because neurons in different layers show different levels of noise correlation (Hansen et al., 2012). The overall level of correlation was comparable across the two subnets (Fig. 7B), making it unlikely that the topography on the array is a simple byproduct of different recording depths across the array.

Shared visual and motor response fields contribute to, but do not fully explain, topography based on dissimilarity

Visuomotor topography has been previously reported for area 8Ar, with central RF's situated ventrally on the prearcuate gyrus and peripheral RF's situated dorsally (Robinson and Fuchs, 1969; Suzuki and Azuma, 1983). Although we were not able to reliably detect this trend within the 4x4 mm patches of cortex beneath our arrays (data not shown), we nevertheless explored the possibility that shared RF location might contribute to our measured dissimilarity matrices (Fig. S9). In the delayed saccade task, we measured the visual and motor response fields of each unit (when possible). Visual responses were calculated as firing rate in a 150 ms window that started 50 ms after target onset, minus a baseline firing rate measured in a 150 ms window immediately before target onset. Peri-saccadic motor responses were calculated as firing rate in a 150 ms window that started

100 ms before saccade initiation, minus a baseline firing rate measured in a 150 ms window starting 400 ms before the saccade. We defined the center of the visual RF of each unit to be the target location that elicited the largest visual response, and the center of the motor RF to be the target location associated with the largest peri-saccadic response. Although our delayed saccades included target eccentricities up to 25 degrees, we were unable to locate the RF center for some units because they did not elicit robust visual or motor responses—presumably, in some cases, because their RFs lay beyond the tested region. For this analysis, we focused only on units that showed significant visual or motor activity relative to the baseline (t-test, $p < 0.05$ for the best target location, Bonferroni corrected). We calculated a matrix of distances between RF centers for all pairs of neurons with well defined RFs in a given experiment, and then calculated the alignment of this matrix with the matrix of response dissimilarities for the same pairs (Fig. S9C). The analysis was performed separately for visual and motor RFs. The RF distances were weakly but significantly correlated with the response dissimilarities. The significance of this correlation suggests that shared RFs contribute modestly to the structure of the subnets. However, the small size of the effect suggests that overlapping RFs, like task-related events (Fig. 4C), explain only a small portion of the total subnet structure.

References

Barnett, L., Buckley, C.L., and Bullock, S. (2009). Neural complexity and structural connectivity. Physical review E, Statistical, nonlinear, and soft matter physics 79, 051914.

Bedard, C., Kroger, H., and Destexhe, A. (2006). Does the 1/f frequency scaling of brain signals reflect self-organized critical states? Physical review letters 97, 118102.

Constantinidis, C., and Goldman-Rakic, P.S. (2002). Correlated discharges among putative pyramidal neurons and interneurons in the primate prefrontal cortex. Journal of neurophysiology 88, 3487-3497.

de la Rocha, J., Doiron, B., Shea-Brown, E., Josic, K., and Reyes, A. (2007). Correlation between neural spike trains increases with firing rate. Nature 448, 802-806.

Galan, R.F. (2008). On How Network Architecture Determines the Dominant Patterns of Spontaneous Neural Activity. PloS one 3.

Green, D.M., and Swets, J.A. (1966). Signal Detection Theory and Psychophysics (New York: Wiley).

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

He, B.J., Zempel, J.M., Snyder, A.Z., and Raichle, M.E. (2010a). The temporal structures and functional significance of scale-free brain activity. Neuron 66, 353-369.

He, B.Y.J., Zempel, J.M., Snyder, A.Z., and Raichle, M.E. (2010b). The Temporal Structures and Functional Significance of Scale-free Brain Activity. Neuron 66, 353-369.

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.

Leopold, D.A., and Logothetis, N.K. (2003). Spatial patterns of spontaneous local field activity in the monkey visual cortex. Reviews in the neurosciences 14, 195-205.

Liu, Y., Hayes, D.N., Nobel, A., and Marron, J.S. (2008). Statistical Significance of Clustering for High-Dimension, Low-Sample Size Data. Journal of the American Statistical Association 103, 1281-1293.

MacQueen, J.C. (1967). Some Methods for classification and Analysis of Multivariate Observations. In Proceedings of 5-th Berkeley Symposium on Mathematical Statistics and Probability (Berkeley, University of California Press), pp. 281-297.

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.

Pernice, V., Staude, B., Cardanobile, S., and Rotter, S. (2012). Recurrent interactions in spiking networks with arbitrary topology. Physical review E, Statistical, nonlinear, and soft matter physics 85, 031916.

Reppas, J.B., and Newsome, W.T. (2008). Decision-outcome encoding in the primate prefrontal cortex. In Society for Neuroscience Meeting (Washington DC).

Rieke, F.M., Warland, D., de Ruyter van Steveninck, R., and Bialek, W. (1997). Spikes: Exploring the Neural Code (Cambridge, MA: MIT Press).

Robinson, D.A., and Fuchs, A.F. (1969). Eye movements evoked by stimulation of frontal eye fields. J Neurophysiol 32, 637-648.

Robinson, P.A. (2012). Interrelating anatomical, effective, and functional brain connectivity using propagators and neural field theory. Physical review E, Statistical, nonlinear, and soft matter physics 85, 011912.

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. (2011). The non-random brain: efficiency, economy, and complex dynamics. Front Comput Neurosci 5, 5.

Sporns, O., Tononi, G., and Edelman, G.M. (2000). Connectivity and complexity: the relationship between neuroanatomy and brain dynamics. Neural networks : the official journal of the International Neural Network Society 13, 909-922.

Suzuki, H., and Azuma, M. (1983). Topographic studies on visual neurons in the dorsolateral prefrontal cortex of the monkey. Experimental brain research Experimentelle Hirnforschung Experimentation cerebrale 53, 47-58.

Teich, M.C., Heneghan, C., Lowen, S.B., Ozaki, T., and Kaplan, E. (1997). Fractal character of the neural spike train in the visual system of the cat. Journal of the Optical Society of America A, Optics, image science, and vision 14, 529-546.

Tsujimoto, S., Genovesio, A., and Wise, S.P. (2009). Monkey orbitofrontal cortex encodes response choices near feedback time. J Neurosci 29, 2569-2574.

Tsujimoto, S., Genovesio, A., and Wise, S.P. (2010). Evaluating self-generated decisions in frontal pole cortex of monkeys. Nat Neurosci 13, 120-126.

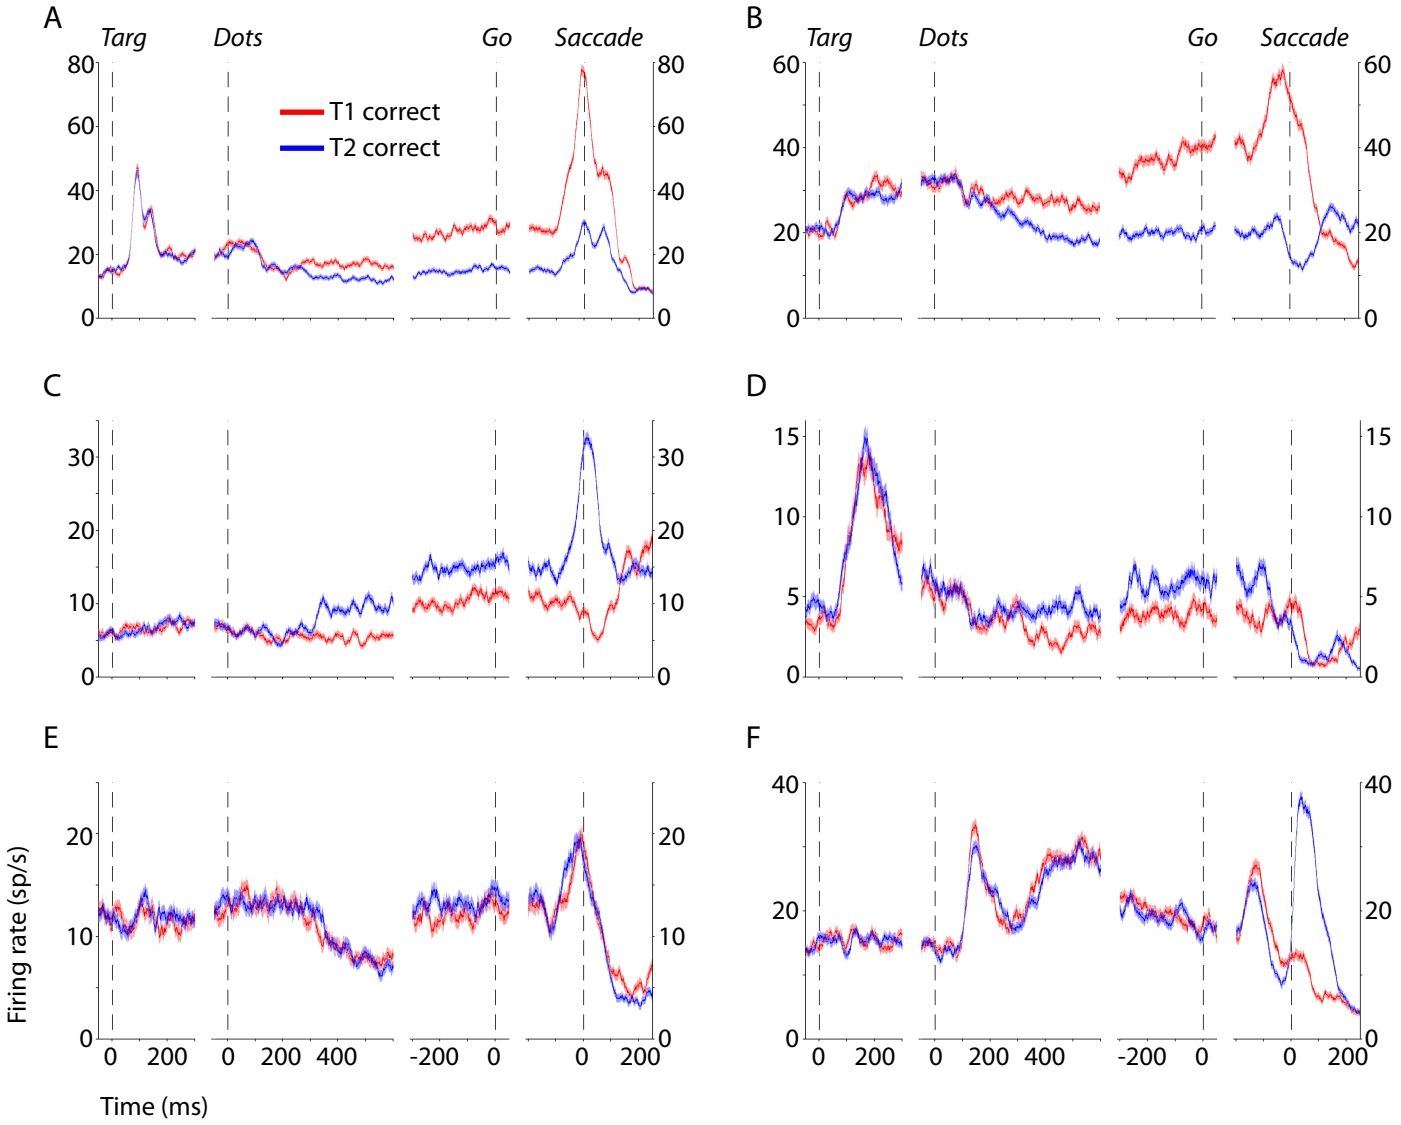

Figure S1, related to Fig. 2 and 3. Prearcuate units exhibit diverse responses during the direction discrimination task. Responses of six example units from the session of Fig. 2C are shown. The responses are aligned to different task events and are grouped based on the chosen target (T1 or T2). **A-B**) Two units that preferred T1 choices and were predictive of the monkey's choice during the motion viewing epoch. Unit A had a sharp phasic visual response to target onset and a sharp motor response prior to and immediately after the operant saccade. Unit B exhibited a more tonic visual response to target onset and was inhibited below baseline on trials in which the monkey ultimately chose T2. **C-D**) Two units that preferred T2 choices and were predictive of the upcoming choice during motion viewing. Unit C lacked visual responses to target onset but responded strongly around the time of the saccade. In contrast, unit D exhibited a strong visual response but lacked a motor response. **E-F**) Two units that were not predictive of the monkey's upcoming choice but, nonetheless, exhibited strong modulation of responses in different task epochs. Unit F is an example of a choice-postdictive unit (see Supplemental Experimental Procedures), which encodes the direction of the recently made operant saccade during the temporal interval between the saccade and receipt (or not) of the reward. The shaded areas represent SEM across trials.

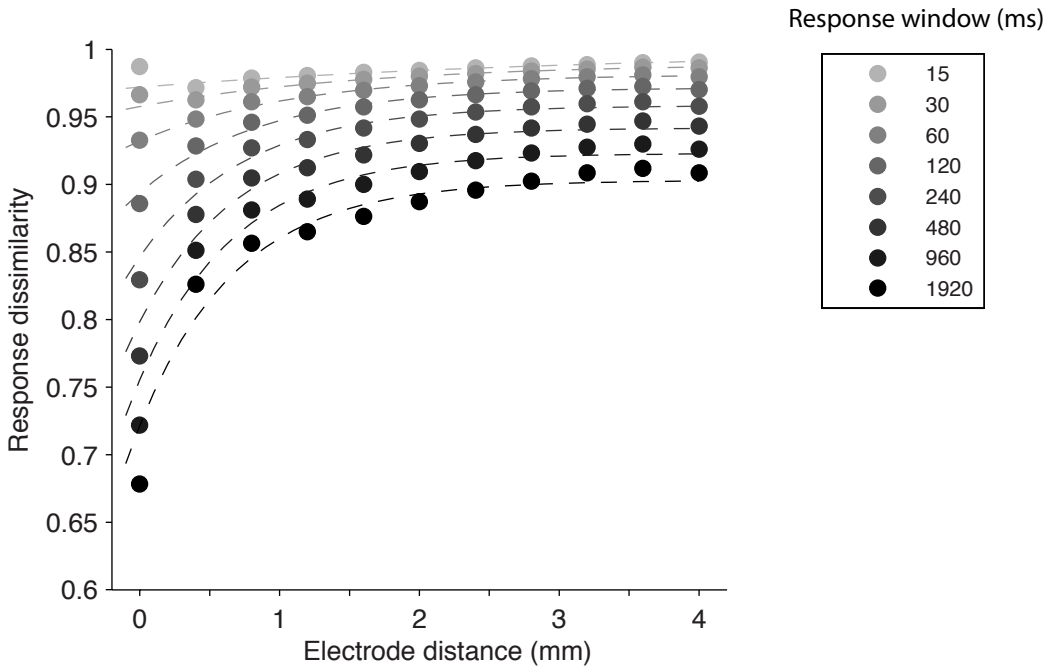

Figure S2, related to Fig. 2. Response dissimilarity varies systematically with the distance between recording electrodes and the size of the time window used for the measurement of responses. Response dissimilarity of a pair of units is defined as one minus the correlation coefficient of their responses across the session (see Experimental Procedures). Dashed lines are fits to the exponential function of Equation S1.

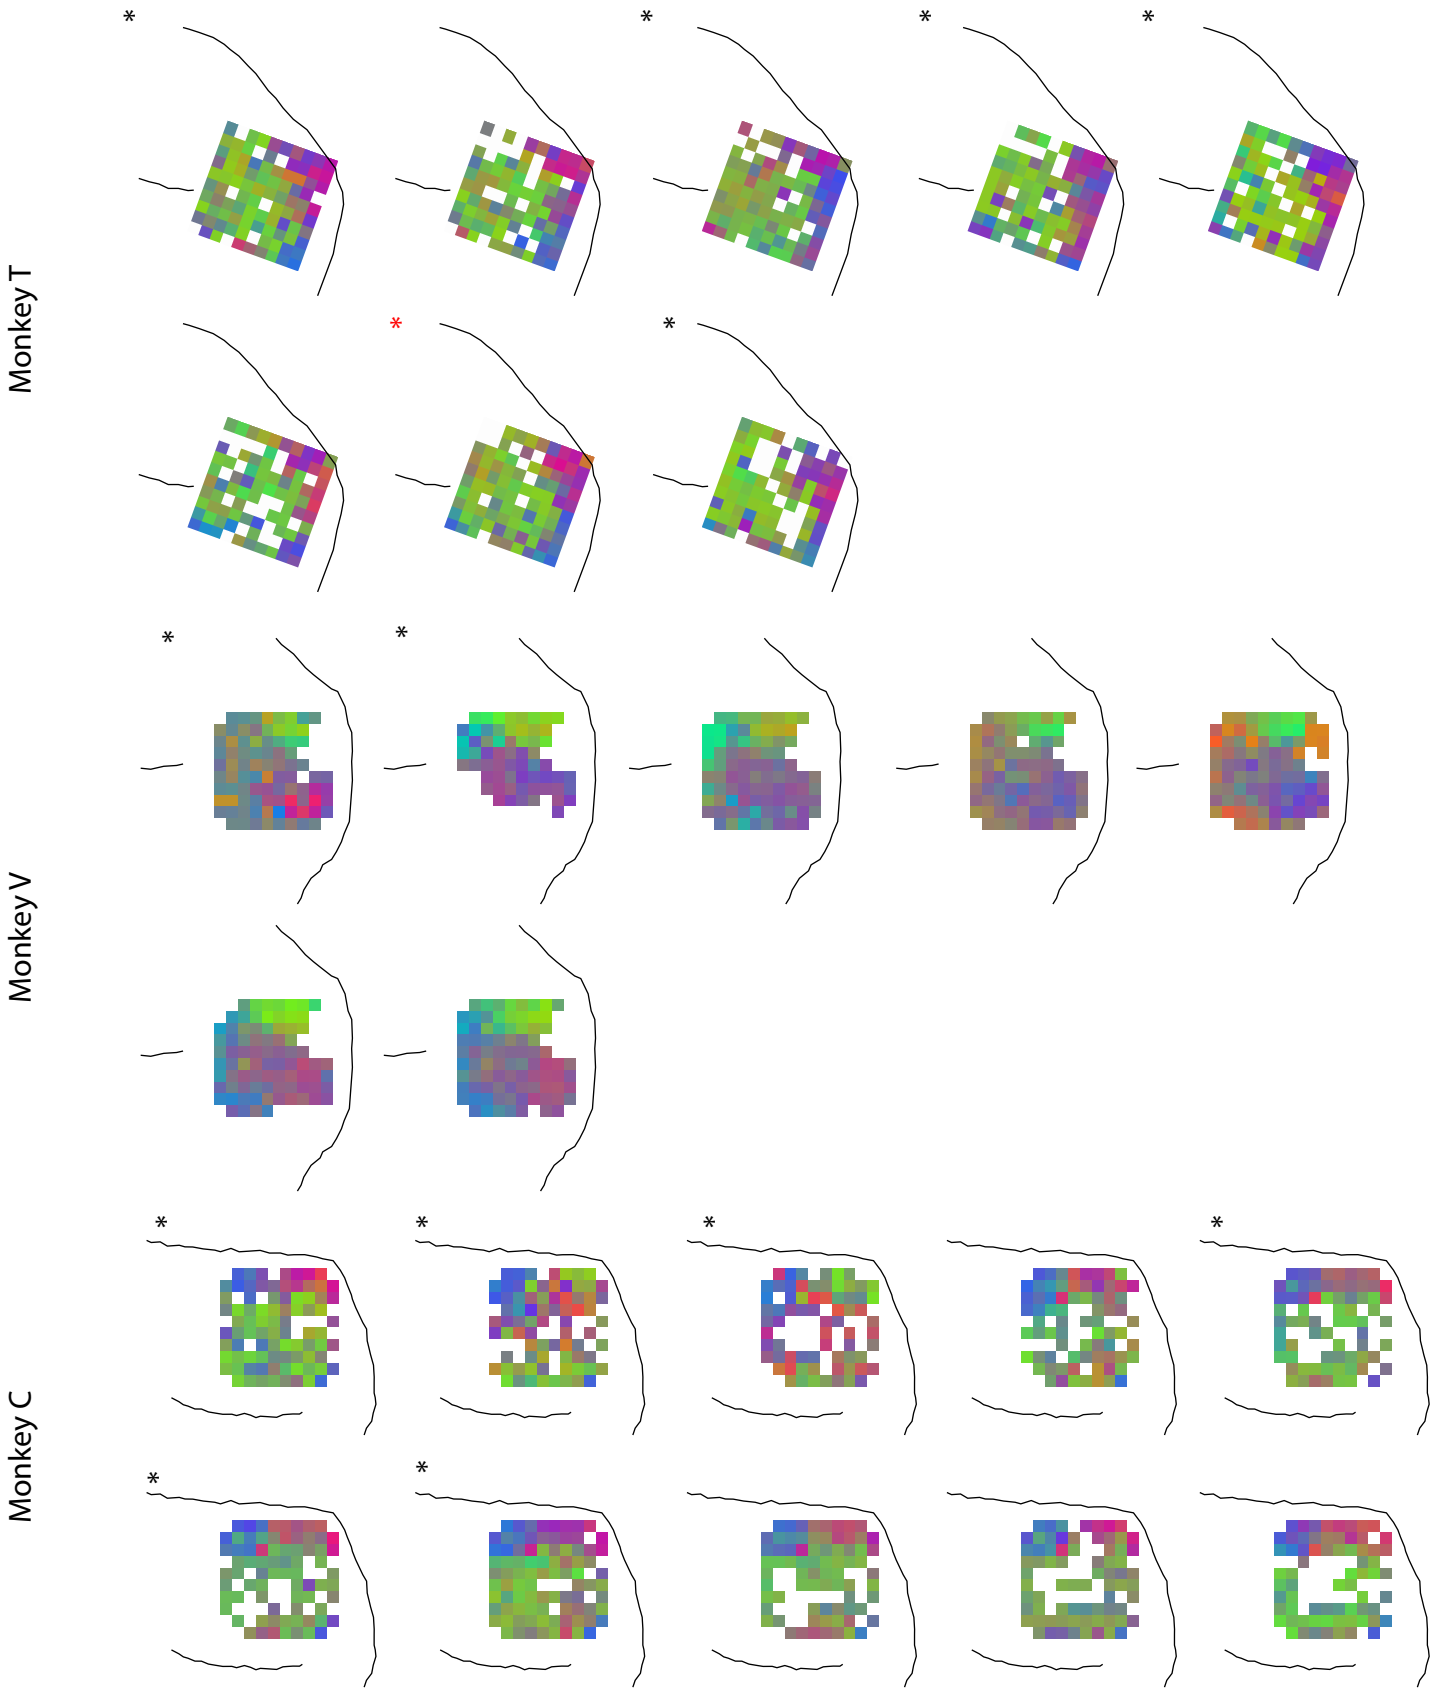

Figure S3, related to Fig. 2. Spatial topography is stable across sessions. All direction-discrimination sessions are depicted for each monkey. Conventions are similar to Fig. 2F-H. The sessions in which the recorded population could not be explained as a single cluster are marked with asterisks (SigClust $p<0.05$). The red asterisk marks the example session that is used in Figures 2C, 3, 4, and 7. Session-to-session similarity of the spatial maps is evident qualitatively, even for sessions in which the clustering was not statistically significant (no asterisks). Precise quantitative comparison of these trends across sessions (using our alignment score) is not feasible because the units on the array frequently differed across sessions, due presumably to slight shifts in the position of the array over time. However, we can calculate a coarser measure—the average dissimilarity across recording sessions of all units recorded by particular pairs of electrodes within sessions. This measure is still suboptimal because the number of electrodes that actually recorded some signal could vary from day to day. Nonetheless, the alignment scores of these electrode-averaged dissimilarities are 0.5±0.1 (mean±s.d.) for monkey T, 0.5±0.2 for monkey V, and 0.4±0.2 for monkey C, indicating strong consistency of the maps across sessions.

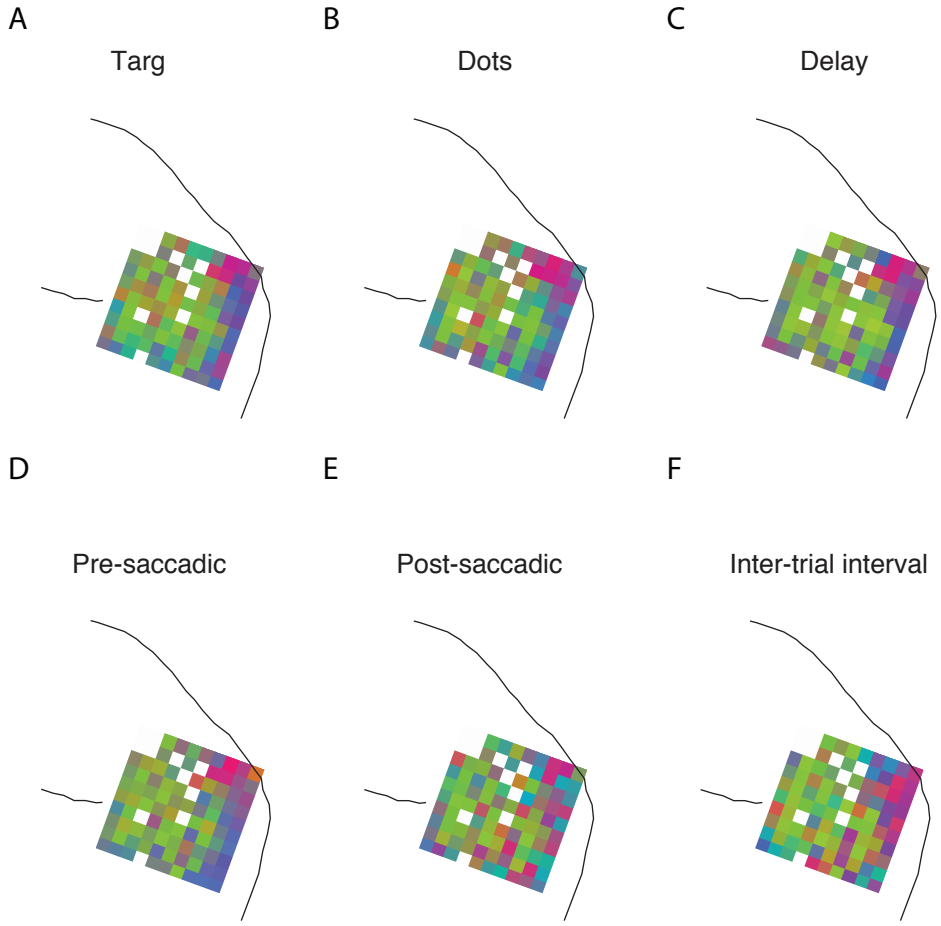

Figure S4, related to Fig. 3. Spatial topography on the array is stable across task epochs. Panels A-F show the array projections of the MDS maps for the same experiment and task epochs illustrated in Fig. 3A-F. However, unlike Fig. 3 where the units inherited their color from the whole session MDS map (Fig. 2C), in this figure the units are assigned new colors based on the circular color map of Fig. 2 and then projected onto the array. The similarity of the array projection with that in Fig. 2F indicates that the topography is largely invariant across task epochs.

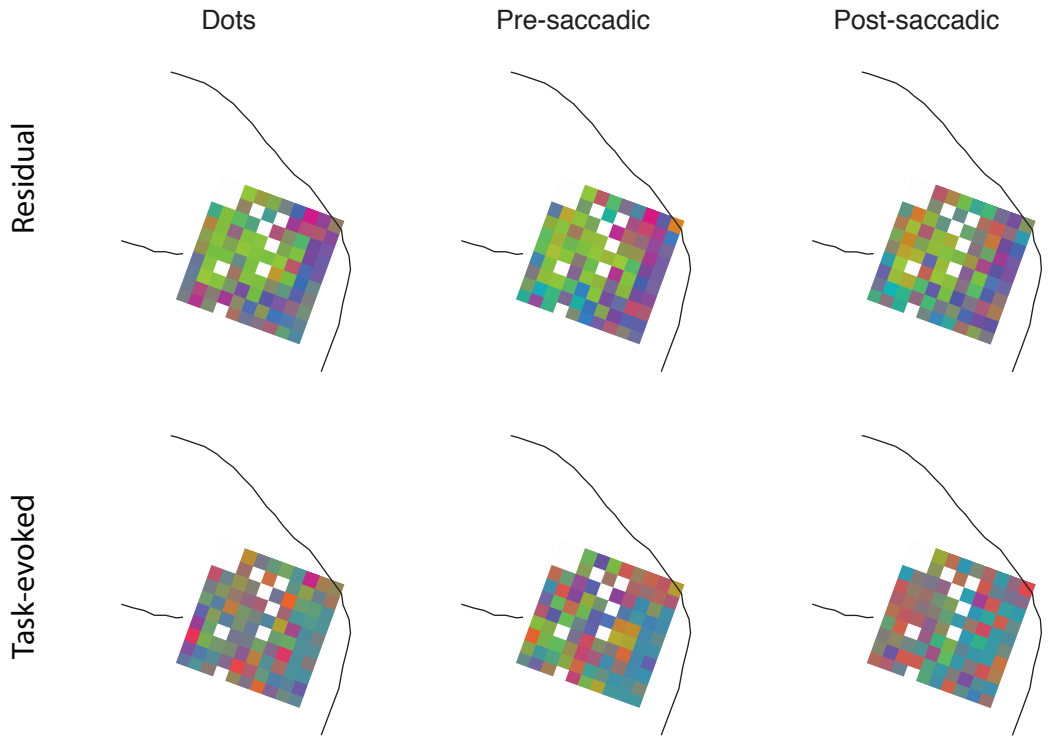

Figure S5, related to Fig. 4. Spatial topography on the array stems from response fluctuations around task-evoked means. The six panels show the array projection of the residual and task-evoked MDS maps in Fig. 4A. The units are recolored according to the circular color map of Fig. 2 before projection on the array. Spatial topography is evident in the residuals (top row) but not in the task-evoked means (bottom row).

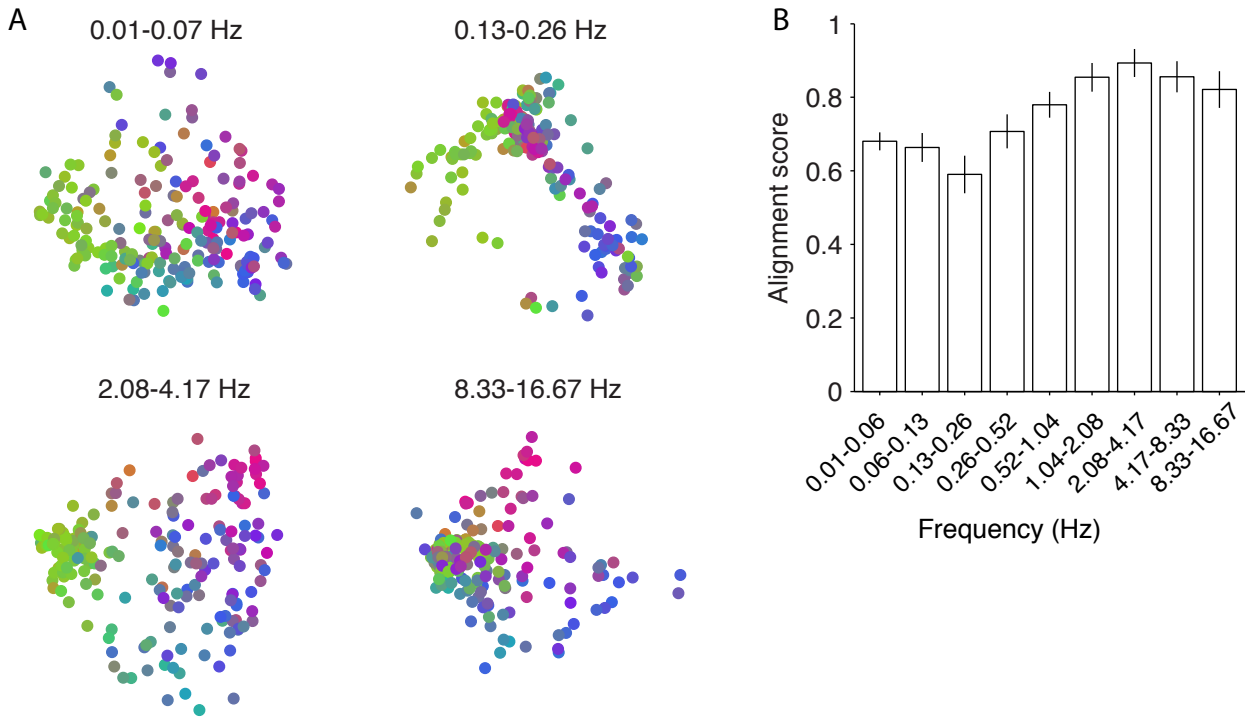

Figure S6, related to temporal frequency analysis. The signal underlying subnet structure is temporally broadband (see the corresponding sections in Results and Supplemental Information). **A**) Two-dimensional MDS plots based on band-pass filtered neural responses for the example session of Fig. 2C. The neural responses across the session were band-pass filtered with nine different frequency bands, four of which are illustrated by MDS plots. The filters spanned three orders of magnitude, ranging from 0.01 Hz to 16.67 Hz, which is the Nyquist frequency for our 30 ms analysis window. **B**) The alignment score of the band-pass filtered response dissimilarities with the whole-session response dissimilarity. The bars show the average alignment across all 25 experimental sessions. Error bars represent 95% confidence intervals.

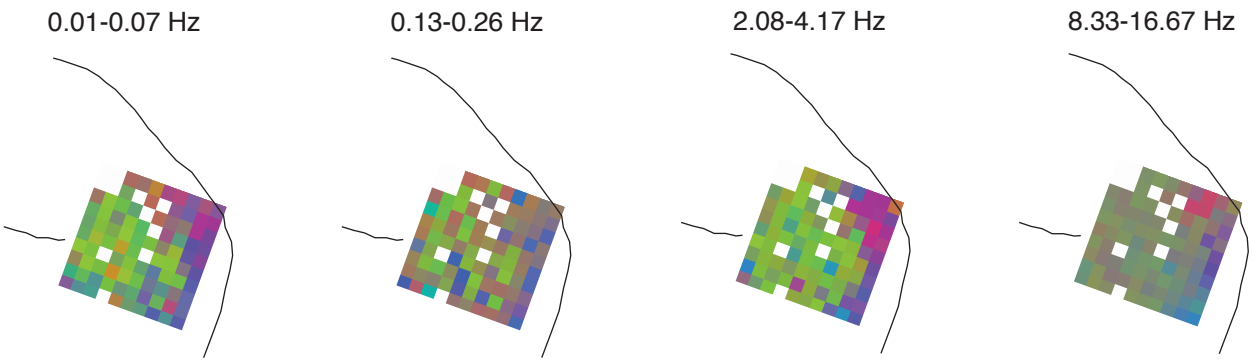

Figure S7, related to temporal frquency analysis. Spatial topography on the array is broadband (see the corresponding sections in Results and Supplemental Information). The four panels show the array projection of band-pass filtered MDS maps in Fig. S6A. The units are recolored according to the circular color map of Fig. 2 before projection on the array.

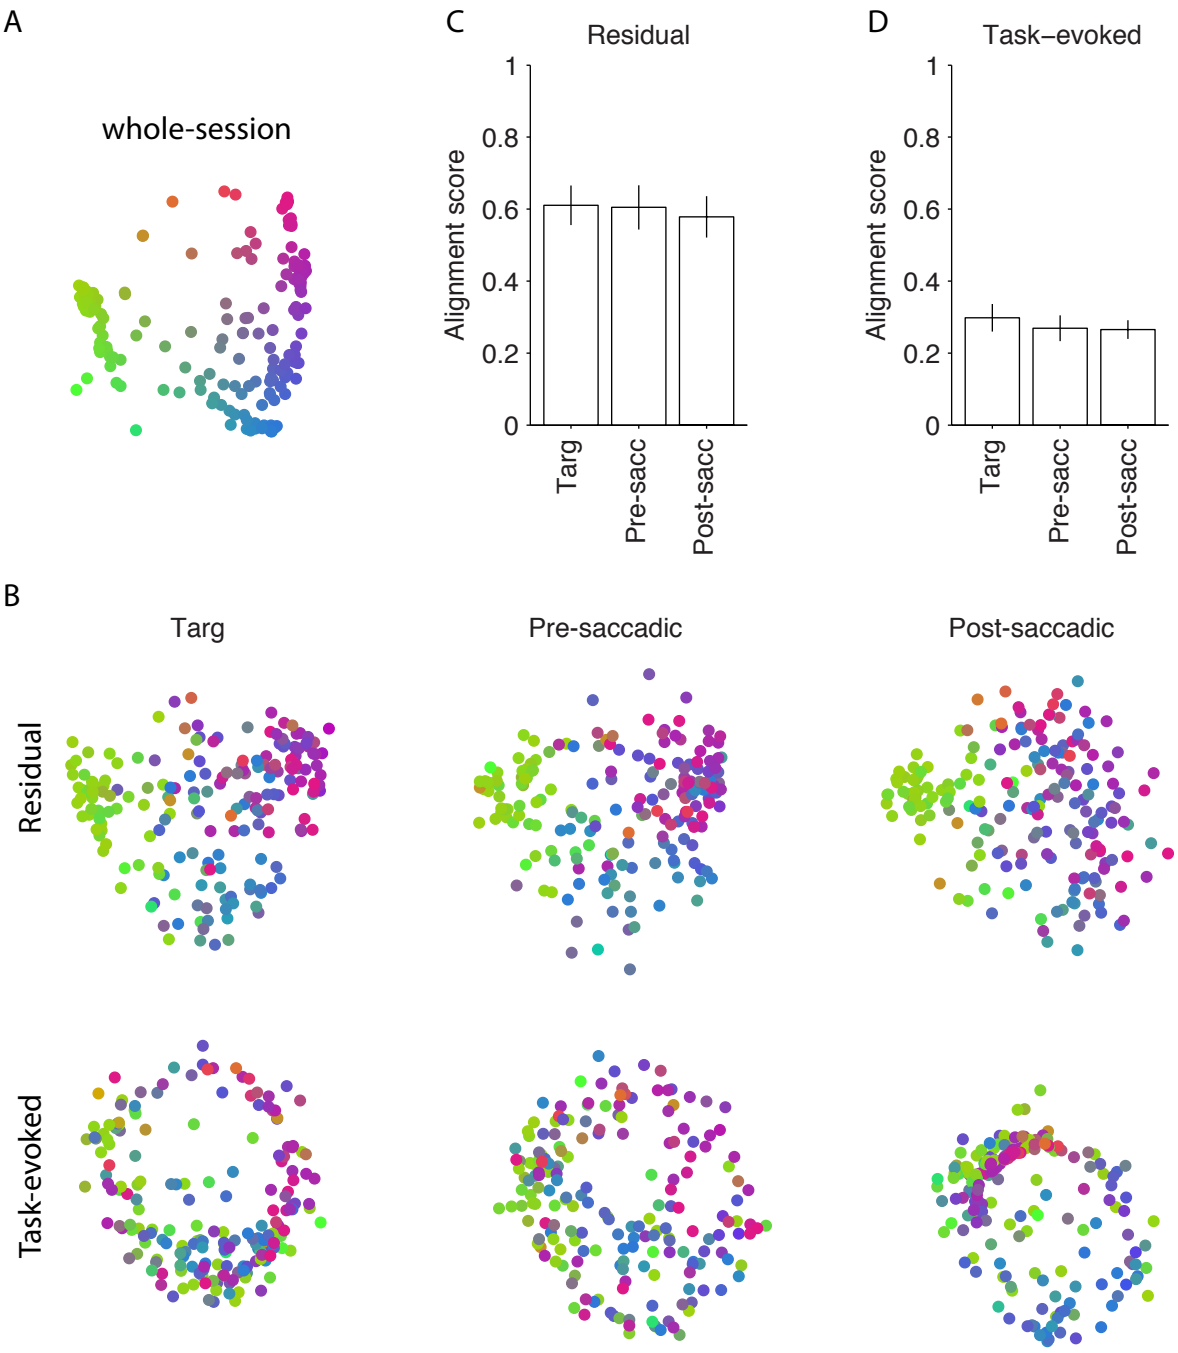

Figure S8, related to Fig. 5. The spatial topography in the delayed saccade task is best explained by common noise rather than task-evoked responses. **A**) Two-dimensional MDS plot of an example session (same as Fig. 5B). **B**) MDS plots based on task-evoked and residual responses in three different task epochs. Conventions are similar to those in Fig. 4A. **C-D**) Alignment of residual and task-evoked response dissimilarities to the whole-session response dissimilarities. Conventions are similar to those in Fig. 4B-C.

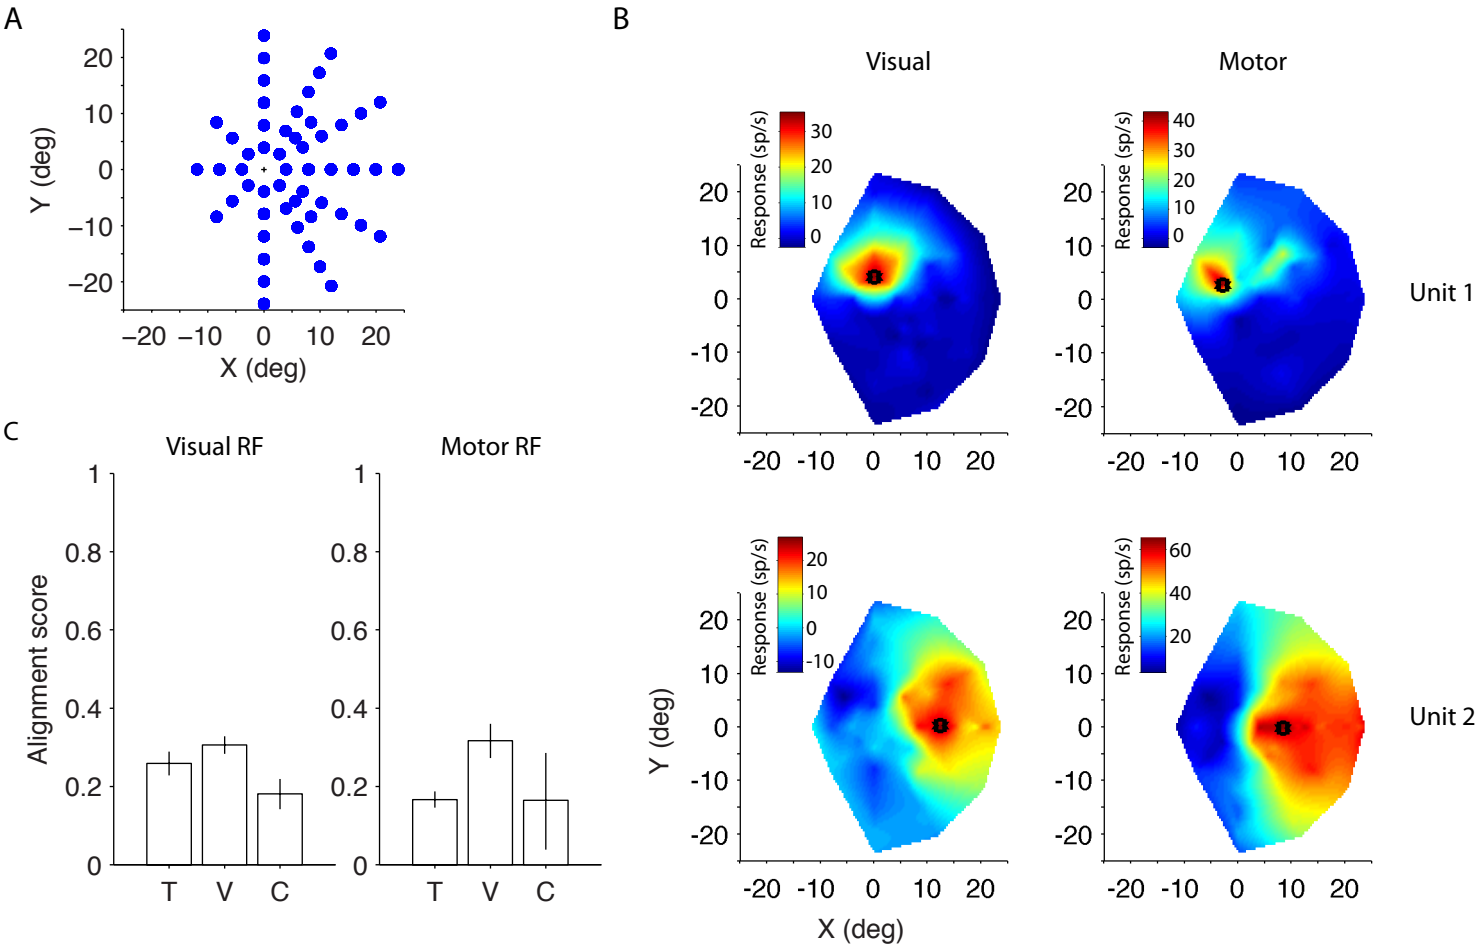

Figure S9, related to Fig. 5. Shared visual and motor response fields contribute to, but do not fully explain, topography based on dissimilarity. **A**) Several target locations (mostly contralateral) were tested in the delayed saccade task, enabling us to measure the response fields of the recorded units. The blue points show the target locations for the example session of Fig. 5. Only one target was shown in each trial. **B**) The magnitudes of visual and motor responses of two example units are depicted for various target locations. Visual responses were calculated as firing rate in a 150 ms window that started 50 ms after target onset, minus a baseline firing rate measured in a 150 ms window immediately before target onset. Peri-saccadic motor responses were calculated as firing rate in a 150 ms window that started 100 ms before saccade initiation, minus a baseline firing rate measured in a 150 ms window starting 400 ms before the saccade. For some target locations the unit responses go below baseline, hence the negative numbers. The RF center, defined as the target location associated with the highest response, is marked with a black asterisk. **C**) The distances between RF centers were calculated to create a RF distance matrix (see Supplemental Information). The alignment of this distance matrix with the whole-session response dissimilarity matrix is shown for the three monkeys. Only units with significant visual or motor responses contributed to the analysis. Error bars represent 95% confidence intervals.

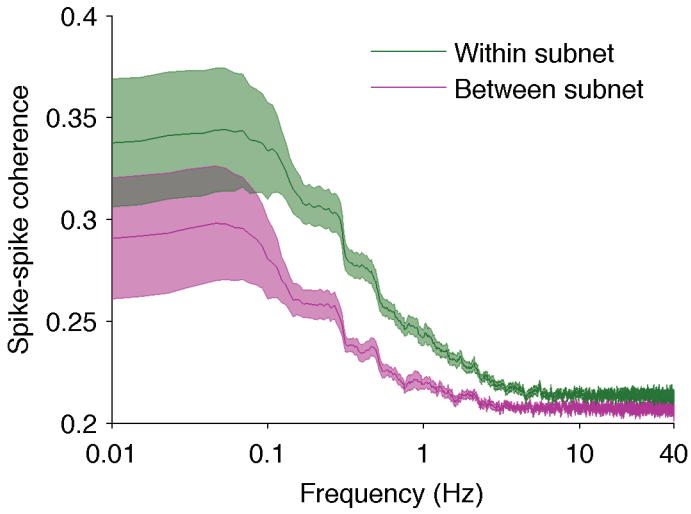

Figure S10, related to Fig. 7. Spike-spike coherence of neighboring units that belong to the same subnet (green) and to different subnets (magenta). Shaded areas indicate standard errors across sessions. The results are compatible with those of Fig. 7 and indicate that subnet identities persist across all epochs and temporal frequencies.

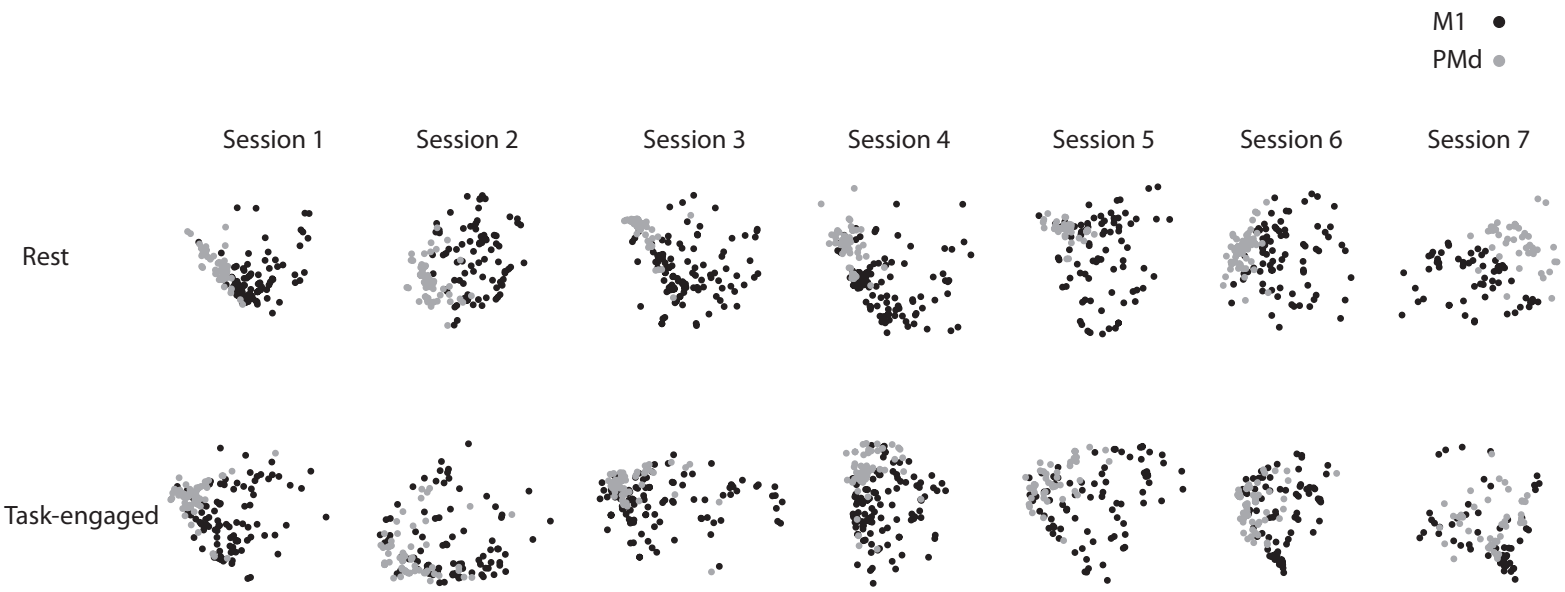

Figure S11, related to Fig. 8. MDS maps of the seven simultaneous recording sessions from the M1 and PMd arrays. M1 and PMd units segregate from each other both during the direction discrimination task (bottom row) and in rest periods between the task-engaged blocks (top row). In the rest periods (15-60 min) the monkey sat calmly in his chair in front of a blank monitor in a semi-dark room. Session 1 is the example depicted in Fig. 8B.