

Mnemonic Encoding and Cortical Organization in Parietal and Prefrontal Cortices

 Nicolas Y. Masse, Jonathan M. Hodnefield, and David J. Freedman

Department of Neurobiology, University of Chicago, Chicago, Illinois 60637

Persistent activity within the frontoparietal network is consistently observed during tasks that require working memory. However, the neural circuit mechanisms underlying persistent neuronal encoding within this network remain unresolved. Here, we ask how neural circuits support persistent activity by examining population recordings from posterior parietal (PPC) and prefrontal (PFC) cortices in two male monkeys that performed spatial and motion direction-based tasks that required working memory. While spatially selective persistent activity was observed in both areas, robust selective persistent activity for motion direction was only observed in PFC. Crucially, we find that this difference between mnemonic encoding in PPC and PFC is associated with the presence of functional clustering: PPC and PFC neurons up to $\sim 700\ \mu\text{m}$ apart preferred similar spatial locations, and PFC neurons up to $\sim 700\ \mu\text{m}$ apart preferred similar motion directions. In contrast, motion-direction tuning similarity between nearby PPC neurons was much weaker and decayed rapidly beyond $\sim 200\ \mu\text{m}$. We also observed a similar association between persistent activity and functional clustering in trained recurrent neural network models embedded with a columnar topology. These results suggest that functional clustering facilitates mnemonic encoding of sensory information.

Key words: macaque monkey; parietal cortex; persistent activity; prefrontal cortex; topological organization; working memory

Significance Statement

Working memory refers to our ability to temporarily store and manipulate information. Numerous studies have observed that, during working memory, neurons in higher cortical areas, such as the parietal and prefrontal cortices, mnemonically encode the remembered stimulus. However, several recent studies have failed to observe mnemonic encoding during working memory, raising the question as to why mnemonic encoding is observed during some, but not all, conditions. In this study, we show that mnemonic encoding occurs when a cortical area is organized such that nearby neurons preferentially respond to the same stimulus. This result provides plausible neuronal conditions that allow for mnemonic encoding, and gives us further understanding of the brain's mechanisms that support working memory.

Introduction

Working memory allows for the temporary storage, and if required, manipulation of sensory information, and is crucial whenever behavioral decisions are required after the relevant stimuli have been extinguished. A neurophysiological hallmark of working memory is the presence of stimulus-specific persistent

activity during the “delay” period following the offset of the to-be-remembered stimulus, which has consistently been observed throughout the frontoparietal network (Funahashi et al., 1989; Colby et al., 1996; Chafee and Goldman-Rakic, 1998; Rainer et al., 1998; Romo et al., 1999; Zaksas and Pasternak, 2006). However, mnemonic encoding in the frontoparietal network does not generalize across all stimulus features or behavioral contexts. For example, the prefrontal cortex (PFC) encodes spatial location and motion direction during the delay period, but a recent study showed that it does not mnemonically encode color during a fine-change detection task (Lara and Wallis, 2014). Similarly, although the posterior parietal cortex (PPC) mnemonically encodes the category membership of visual motion stimuli as a result of categorization training (Freedman and Assad, 2006; Swaminathan and Freedman, 2012), mnemonic encoding of motion direction during a delayed-matching task (before categorization training) is weak (Sarma et al., 2016), despite robust mnemonic encoding in the upstream, motion-selective medial

Received Dec. 21, 2016; revised March 31, 2017; accepted May 2, 2017.

Author contributions: N.Y.M. and D.J.F. designed research; N.Y.M. and J.M.H. performed research; N.Y.M. analyzed data; N.Y.M. and D.J.F. wrote the paper.

This work was supported by National Institutes of Health R01EY019041 and R01MH092927, and National Science Foundation Career Award NCS 1631571. We thank the staff at the University of Chicago Animal Resource Center for expert assistance.

The authors declare no competing financial interests.

Correspondence should be addressed to either Dr. Nicolas Y. Masse or Dr. David J. Freedman, Department of Neurobiology, University of Chicago, 5812 S. Ellis Avenue, MC0912, P-419, Chicago, IL 60637. E-mail: masse@uchicago.edu or dfreedman@uchicago.edu.

DOI:10.1523/JNEUROSCI.3903-16.2017

Copyright © 2017 the authors 0270-6474/17/376098-15\$15.00/0

superior temporal area (Mendoza-Halliday et al., 2014). This raises the question of why frontoparietal stimulus-selective persistent activity exists for some, but not all, visual features.

On a different level, several modeling studies have examined the possible circuit mechanisms underlying persistent activity. It is thought that recurrent excitation within local networks, subserved by the slow kinetics of NMDA receptors, can allow for stimulus-selective activity to persist after the stimulus is removed (Wang, 1999; Wang et al., 2013). Although not explicitly stated, these models suggest a level of cortical organization, in which neurons selective for similar stimuli are anatomically clustered. Recurrent excitation between neurons is enhanced when all neurons prefer similar stimuli (or equally valid, recurrent excitation will cause all neurons within the group to selectively respond to similar stimuli). As most synaptic connections are between neurons separated from one another by less than several hundred microns (Perin et al., 2011; Levy and Reyes, 2012), the implication is that persistent activity is associated with the presence of functional clustering. In agreement with this notion, areas such as PPC and PFC that mnemonically encode spatial location (Colby et al., 1996; Chafee and Goldman-Rakic, 1998; Rainer et al., 1998), also contain retinotopic maps (Serenio et al., 2001; Silver and Kastner, 2009; Patel et al., 2010). However, if we are to understand why stimulus-specific persistent activity generalizes to some, but not all, visual features, what is needed are studies that directly compare how different cortical areas encode different features in working memory.

In this study, we ask whether persistent activity in the PPC and the PFC is associated with the presence of functional clustering. We take advantage of a semichronic recording system, which allows us to track the relative anatomical location of our neural recordings across sessions, allowing us to measure the similarity in stimulus tuning between spatially clustered neurons. We find that neurons whose tuning is similar to that of its neighbors are more likely to mnemonically encode the stimulus during the delay period of the task. Furthermore, we find that this relation between local tuning similarity and mnemonic encoding occurs in recurrent neural network models, embedded with a columnar topology, which have been trained to perform a delayed-matching task. Finally, we find that the local field potentials (LFPs), whose activity is thought to reflect the sum of synaptic currents within the local volume (Buzsáki et al., 2012), are spatially selective in both PPC and PFC within 70 ms of stimulus onset, whereas motion selectivity was much weaker. These results suggest that the presence of persistent activity depends upon functional clustering, and that inputs conveying spatial selectivity into PPC and PFC, but not motion inputs, are themselves functionally organized.

Materials and Methods

Behavioral tasks and display. Two male monkeys (*Macaca mulatta*) were trained to perform a spatial-based and two different motion-based tasks. Stimulus presentation, task events, rewards, and behavioral data acquisition for both tasks were accomplished using MonkeyLogic software (<http://www.brown.edu/Research/monkeylogic>) running in MATLAB (The MathWorks) (Asaad et al., 2013). The spatial-based task was a delayed memory saccade task. The monkeys had to maintain fixation (within 2°) on a central point for 500 ms, followed by a visual target presentation in one of eight locations for 307 ms, and finally by a delay period of 1013 ms. The central fixation point was then extinguished, and the monkey had to saccade to the location of the remembered visual target to receive a reward. The angular direction of the visual target was a multiple of 45° above horizontal, and its eccentricity was 7°. We refer to this task as the spatial task.

The two motion-based tasks were a delayed match-to-sample (DMS) and a delayed match-to-category (DMC) task. The DMS and DMC tasks have also been previously described (Sarma et al., 2016). In both tasks, the monkeys had to maintain fixation (within 2.2°) on a central point for 500 ms, followed by a 667 ms sample motion stimulus, followed by a 1013 ms delay, and then a 667 ms test motion stimulus. Monkeys released a manual lever to indicate whether a test stimulus was the same direction (DMS) or same category (DMC) as the previously presented sample. The motion stimulus consisted of circular patches of high-contrast, 100% coherent random dots displayed at a frame rate of 75 Hz. The motion stimuli had a diameter of 6.0° and dots moved at 12°/s. We refer to these two tasks as the motion task.

Both monkeys were highly trained on the DMS task before recording (>150 sessions each), and we recorded PPC and PFC activity during 7 DMS sessions, after which, the 2 monkeys learned, through trial and error, to perform the DMC task while we continued to record from PPC and PFC. We recorded for an additional 26 DMC sessions for Monkey Q and 30 DMC sessions for Monkey W. Motion direction selectivity was not appreciably different between the DMS and DMC sessions (data not shown); we thus combined all sessions for the analysis in this study. Although we have previously shown that extensive categorization training (beyond that experienced by the monkeys in this study) can alter delay period category representations in lateral intraparietal (LIP) area (Sarma et al., 2016), we found that only including DMS sessions or only DMC sessions does not qualitatively change our main results.

Gaze positions were measured and recorded at a sampling rate of 1 kHz using an EyeLink 1000 optical eye tracker (SR Research). Visual stimuli were presented on a 21 inch color CRT monitor (1280 × 1024 resolution, 57 cm viewing distance).

Electrophysiological recording. We used two 32-channel semichronic recording microdrives (Gray Matter Research) to record from PPC and PFC. MRI scans were used to guide chamber placement. For PPC recordings, chambers were placed over the intraparietal sulcus, ~2.0 mm posterior to the intra-aural line and ~14.0 mm lateral from the midline for Monkey Q, and ~2.0 mm anterior to the intra-aural line and ~13.0 mm lateral from the midline for Monkey W. For PFC recordings, chambers were placed over the principal sulcus, ~29.0 mm anterior to the intra-aural line and ~20.0 mm lateral from the midline for Monkey Q, and ~33.0 mm anterior to the intra-aural line and ~22.0 mm lateral from the midline for Monkey W. Each microdrive system contained 32, 125 μ m tungsten microelectrodes (Alpha-Omega). Adjacent electrodes were spaced 1.5 mm apart.

Before each session, we advanced electrodes (i.e., lowered them into the brain) by between 0 and ~1 mm to optimally record the spiking activity of well-isolated neurons. For PFC recordings, we recorded from neurons as soon as we entered cortex. For PPC recordings, we wanted to target the LIP area. Thus, we advanced electrodes below the intraparietal sulcus, which was inferred based on its cortical location and a lack of spiking activity within the sulcus. We recorded from all neurons for which we could reliably sort action potentials.

Neural data were collected using a Plexon multichannel acquisition processor data acquisition system. Action potential data were sampled at 40 kHz and high-pass filtered at 250 Hz with a second-order Butterworth filter, and low-pass filtered at 8000 Hz with a third-order Butterworth filter. LFPs were sampled at 1000 Hz and filtered between 0.7 and 300 Hz using a first-order Butterworth filter. We removed line noise from the LFPs by filtering the signals with a symmetric third-order Butterworth filter with corners at 59 and 61 Hz.

All surgical and experimental procedures followed the University of Chicago's Animal Care and Use Committee and National Institutes of Health guidelines. Monkeys were housed in individual cages under a 12 h light/dark cycle. Behavioral training and experimental recordings were conducted during the light portion of the cycle. Neurophysiological signals were amplified, digitized, and stored for offline spike sorting to verify the quality and stability of neuronal isolations.

Neuron selection. Because it is difficult to accurately measure stimulus tuning with low spike counts, we only included neurons with mean spike rates (measured between stimulus onset and the end of the delay) >1 Hz for both the spatial and motion tasks. Additionally, we only included

neurons that had at least 8 valid trials for each of the 8 spatial directions during the spatial task, and 8 valid trials for each of the 6 motion directions during the motion tasks. A minimum of 8 trials was chosen so that we could devote 75% of trials for training the linear classifiers with at least two trials per direction for testing the classifier (described below).

Our semichronic recording system allowed us to record from what was potentially the same neuron across multiple recording sessions. However, for this study, we did not want the same neuron recorded across multiple sessions to contribute more than once in any of our analyses. Thus, we eliminated all neurons that we potentially recorded from on a previous session. To determine whether the same neuron was recorded across different sessions, we first found pairs of neurons that were recorded from the same electrode on sequential recording sessions in which the electrode depth was adjusted $<62.5 \mu\text{m}$ (half a rotation of the screw driven actuator that controls electrode depth) between sessions. We note that we only adjusted the electrode depth at the start of the session if no isolated neuron was present. We then eliminated the neuron from the earlier session if the waveform width (see below) differed by no more than 25 μs . So, for example, if the same neuron was recorded across 10 consecutive recording sessions, we would use the last session of data for analysis, and discard the data from the first nine sessions.

To calculate to waveform width (time from trough to peak), we first interpolated the waveform of each action potential to have 1 μs resolution. Next, we aligned all action potentials using the time of their global minimum. We discarded all waveforms in which there existed a local minimum or maximum between the global minimum (trough) and the global maximum (peak) of the waveform. We then computed the mean of these aligned waveforms and measured the time between the trough and the peak.

At the start of each session, we usually either kept the electrode at the same depth or advanced the electrode to isolate a new neuron. However, electrodes could only be advanced a total of $\sim 19 \text{ mm}$, with several millimeters required to reach cortex. Thus, in cases where an electrode reached its maximum depth, we slowly began to retract it (usually by 0 to $\sim 2 \text{ mm}$ per session) to attempt to locate new neurons. This occurred more frequently for our PPC recordings because electrodes were initially lowered below the intraparietal sulcus (within several millimeters of the maximum depth for some PPC channels) before we started collecting data, whereas we collected neural data from PFC as soon as we encountered cortex. To ensure that this difference did not bias our comparison between PPC and PFC, we excluded all neurons which we recorded after we began raising the electrode.

Single-neuron selectivity. To quantify neuronal spatial location and motion direction selectivity (see Fig. 3C,D), we wanted to measure the percentage of variance in the spike count that is explained by the stimulus. However, the traditional measure of percentage of explained variance (PEV) is positively biased for small sample sizes. Thus, we calculated each neuron's spatial and motion direction selectivity using a normalized PEV (Buschman et al., 2011) as follows:

$$\text{normalized PEV} = \frac{\text{SS}_{\text{between groups}} - \text{df} \times \text{MSE}}{\text{SS}_{\text{total}} + \text{MSE}}$$

where $\text{SS}_{\text{between groups}}$ is sum of squares between groups, SS_{total} is the total sum of square, MSE is the mean squared error, and df is the number of degrees of freedom. This normalized metric avoids the positive bias in the traditional measure.

Population decoding. Similar to our previous studies (Swaminathan et al., 2013; Sarma et al., 2016), we also measured stimulus selectivity across our neuronal (see Fig. 3A,B) and LFP (see Fig. 7C) populations by measuring how accurately we could decode the motion direction or spatial location using multiclass support vector machines (SVMs). In this approach, we trained SVM classifiers, using a linear kernel, on neural data to decode the motion direction or spatial location for different trials, and then measured decoder accuracy by comparing the motion direction or spatial location predicted by the classifier with the actual motion direction or spatial location. Specifically, we defined the decoder accuracy as the mean dot product between the spatial location or motion direction predicted from the decoder and the actual spatial or motion direction.

Thus, a score of 1 indicates perfect decoding, 0 indicates chance decoding, and a score of -1 indicates that the decoded direction is opposite the actual direction.

When decoding the motion direction or spatial location from the spiking activity of neurons, we summed spike counts using a causal 200 ms boxcar filter; and when decoding motion direction or spatial location from the LFPs, we causally filtered LFPs using a 10 ms boxcar filter. Spike counts and filtered LFPs were both normalized between -1 and 1 for each neuron or channel and for each time point.

As the number of neurons per session varied, we decided to decode the spiking activity from a surrogate population of neurons. Thus, trials used for training or testing the decoder did not necessarily contain spiking activity that occurred simultaneously across neurons, but rather includes the spike count of neurons in response to the same spatial location or motion direction, but likely recorded at different times. In contrast, because the number of LFP channels remained constant across sessions, trials used for training or testing the decoder always used LFP data that were recorded simultaneously. Thus, trials used for training or testing the decoder always contained LFPs recorded at the same time across all electrodes. We calculated the decoder accuracy for the spatial and motion tasks for each individual session before averaging our results across all sessions to obtain the mean LFP decoder accuracy.

To measure whether decoding accuracies were significantly different between the PPC and PFC populations, we needed to generate distributions of decoding accuracies for both populations. To do so, we used a bootstrap procedure in which we randomly selected, with replacement, both neurons (or LFP channels) and trials for both training and testing the decoder. Given that decoding accuracy will depend both on which neurons and which trials are included, randomly selecting both fully captures the variance in our model.

To ensure a fair comparison when comparing PPC and PFC decoding accuracies, we wanted to include an equal number of neurons to train and test each classifier. Thus, if there were n_1 PPC neurons, n_2 PFC neurons, and, for example, $n_1 < n_2$, we randomly sampled with replacement n_1 neurons from each population to train the classifier. We calculated the decoding accuracy from PPC and PFC for each monkey individually before averaging the results across the 2 monkeys.

We randomly sampled trials using fourfold cross-validation, in which we randomly selected 75% of trials for training the decoder and the remaining 25% for testing the decoder. For each of the six motion directions and eight spatial locations, we randomly sampled, with replacement, 20 trials from each neuron to train the decoder (from the 75% of trials set aside for training), and 20 trials to test the decoder (from the 25% of trials set aside for testing). We repeated this procedure a total of four times, in which we alternated which 25% of trials were set aside for testing, and calculated the mean of the four decoder accuracies at each time point.

We then repeated this procedure of sampling neurons and trials 100 times to create a decoder accuracy distribution for each time point. To determine whether the decoder accuracy was significantly different between PPC and PFC at any time point, we first calculated the 10,000 (100×100) differences between all pairs of decoder accuracies. The difference was deemed significant if $>99.5\%$ of decoder accuracies from one area were greater than the other (equivalent to $p < 0.01$ for a two-sided test).

Measuring functional clustering. We measured the strength of the spatial location (see Fig. 4A) and motion direction (see Fig. 4B) functional clustering by calculating the similarity between the preferred spatial locations or motion directions of nearby neurons. First, we modeled the spike count of neuron i at time t , $z_i(t)$, as a linear function of the spatial location or motion direction, d as follows:

$$z_i(t) = \text{baseline} + dH_i(t) + \varepsilon_i(t),$$

where $\varepsilon_i(t)$ is a Gaussian noise term and the vector $H_i(t)$ relates the stimulus direction to the spike count. For the sake of clarity, for any trial, $z_i(t)$ is a scalar giving the spike count at time t , d is a 1×2 unit vector indicating the spatial location or motion direction for the trial, and $H_i(t)$ is a 2×1 vector that relates the stimulus to the spiking activity.

The angle of $H_i(t)$ is the preferred direction of the neuron at time t , and its magnitude indicates the expected increase in the spike count from baseline when the stimulus matches the preferred direction of the neuron. Thus, the preferred direction of a neuron, represented as a unit vector, is as follows:

$$PD_i(t) = \frac{H_i(t)}{|H_i(t)|}$$

We can calculate how well this linear model fit the data for each neuron i and time point t , indicated by $w_i(t)$, by comparing the variance in the residuals with the variance in the spike counts as follows:

$$w_i(t) = 1 - \frac{\text{var}(\hat{z}_i(t) - z_i(t))}{\text{var}(z_i(t))}$$

Where the fitted spike count based on the model is as follows:

$$\hat{z}_i(t) = \text{baseline} + dH_i(t)$$

Next, we found all pairs of neurons recorded on the same electrode but during different sessions that were located between 62.5 and 750 μm apart (corresponding to 0.5–6 rotations of the screw-driven actuator that controls electrode depth). We then calculated the similarity between each pair of nearby neurons, $s_{i,j}(t)$, as the dot-product between their preferred directions, weighted by the geometric mean of their linear model fits as follows:

$$s_{i,j}(t) = \sqrt{w_i(t)w_j(t)}PD_i^T(t)PD_j(t).$$

We then summed the similarity scores for all pairs of nearby neurons within each cortical area, and divided by the sum of the geometric means of their respective model fits as follows:

$$\text{similarity}(t) = \frac{\sum_{i,j} s_{i,j}(t)}{\sum_{i,j} \sqrt{w_i(t)w_j(t)}}$$

This was calculated for PPC and PFC in each monkey individually, before averaging the results across the 2 monkeys.

For Figure 4C, D, we wanted to calculate the correlation between the tuning similarity and the spatial and motion selectivity across time. First, for each neuron (referred to as neuron A), we found all other neurons recorded from the same electrode during different sessions that were located from 62.5 to 750 μm away. We discarded neuron A if there were not at least two other neurons within this range. As an example, suppose there were three neurons that satisfied these criteria, referred to as neurons B–D. We would first calculate (1) the tuning similarity between B and C, B and D, C and D (as described above) and take the mean, and then (2) calculate the linear model fit of neuron A (given by $w_i(t)$). We calculated these values for all neurons within each cortical of each individual monkey for both tasks, and then calculated the correlation between the tuning similarity values and the model fit values.

Recurrent neural network model. In Figure 6, we trained recurrent neural networks using the Pycog framework (Song et al., 2016) to perform a variant of the delayed motion direction matching task performed by the 2 monkeys. The model (see Fig. 6A) consisted of 72 motion direction-selective neurons, in which neuron i 's response, r_i , to motion stimulus with direction θ was as follows:

$$r_i = \alpha + \alpha \cos(\theta - \theta_i)$$

where the preferred direction of neuron i , θ_i , varied in 5 degree intervals such that the 72 neurons uniformly covered the 360 degrees. We set α to 2.5, although our results were insensitive to the exact value. These 72 neurons projected onto 180 recurrently connected neurons with a connection probability of 25%. The 180 neurons in recurrent network were organized into 6 “columns” of 24 excitatory and 6 inhibitory neurons. Neurons within a column were connected with 50% probability, and neurons between columns were connected with 5% probability. These 180 neurons projected onto two output neurons, a “match” and a “non-match” neuron, with 50% connection probability. All connection

weights to and from the recurrent network were drawn from a uniform distribution; connection weights within the recurrent network were drawn from a gamma distribution such that the sum of excitatory and inhibitory weights was equal. As above, our results were insensitive to the exact values chosen for the connection probabilities.

The network was trained so that the “match” neural response increased from 0.2 to 1.0 (values arbitrarily set) for a matching test stimulus, and that the “nonmatch” neural response increased from 0.2 to 1.0 for a nonmatching test stimulus. The values 0.2 and 1.0 were arbitrarily chosen, and our model was insensitive to exact values provided there was a sufficient gap between the two values. To ensure that our results were not specific to exact connection weights, we randomly initialized and trained 20 different networks, and analyzed the results of all 20 networks.

Experimental design and statistical analysis. The data in this study consist of 343 PPC and 588 PFC neurons recorded from two male macaque monkeys. Most of our statistical analysis involved measurements that were not from single neurons per se, but from the PPC or PFC populations as a whole. This included our population decoding analysis (see Figs. 3A, B, 5A, F, 9C), our measurement of tuning similarity between nearby neurons (Figs. 4A, B, 5B, G, 6A–D, 8E), and the correlation between tuning similarity and stimulus selectivity (Figs. 4C, D, 5C, D, H–J, 8F, G). For these population-based measurements, we used a bootstrap approach to perform statistical comparisons. The bootstrap approach for the population decoding analysis was slightly different from the others as it involved randomly sampling both neurons and trials (described in Population decoding).

For the tuning similarity measurement, if there were n pairs of nearby neurons in the population, we would randomly sample, with replacement, n pairs, and then calculate the weighted tuning similarity (as described in Measuring functional clustering) from this group of neurons. Because we wanted to equally weigh the contribution of each monkey toward this score, we performed this process for each individual monkey and took the average of the two values. We then repeated this entire process 1000 times to generate a distribution of tuning similarity values.

The bootstrap approach for the correlation between tuning similarity and stimulus selectivity was similar. If there were n groups of nearby neurons, we would randomly sample, with replacement, n groups from this population, and calculate the tuning similarity and the mean stimulus selectivity for each of these n groups (as described in Measuring functional clustering). We would then calculate the Pearson correlation coefficient between these two values from the n groups. As above, we repeated this process for each individual monkey and took the average of the two correlation values. We would repeat this entire process 1000 times to generate a distribution of tuning similarity and stimulus selectivity correlation values.

We also used a bootstrap approach for our single-cell measurement of stimulus selectivity (see Fig. 3C, D). Although we could have used a t test or a similar test to perform our statistical analysis, we wanted to equally weigh the contribution of each monkey, which was simplified by using a bootstrap approach, in which distributions for each monkey were calculated individually before being combined. Specifically, if there were n neurons in the population, we would randomly sample, with replacement, n neurons, and then calculate the normalized PEV (as described in Single-cell selectivity) of this group of neurons. We repeated this process for each individual monkey and took the average of the two normalized PEV values. We would then repeat this entire process 1000 times to generate a distribution of PEV values.

We deemed a measurement significantly >0 if 99.5% or more of the bootstrapped values were >0 (equivalent to a $p < 0.01$ for a two-sided test). To determine whether two bootstrapped distributions were significantly different from each other, we computed all (1000×1000) pairwise differences, and deemed the difference significant if $\geq 99.5\%$ ($\leq 0.5\%$) of the pairwise differences were >0 .

All other statistical tests used in this study are described in Results.

Results

The goal of this study was to examine the relationship between mnemonic encoding in the frontoparietal network and the presence of functional clustering (i.e., nearby neurons preferentially

responding to similar stimuli). We used two stimulus features, spatial location and motion direction, which we examined using different behavioral tasks. We measured spatial location selectivity using a delayed memory saccade task (Fig. 1A; see Materials and Methods), in which the monkeys had to remember the spatial location of a visual target that was flashed in one of eight spatial locations. We refer to this task as the spatial task. Motion direction selectivity was measured using DMS or DMC tasks (Sarma et al., 2016) (Fig. 1B; see Materials and Methods), in which monkeys had to determine whether the motion direction of sample and test stimuli, separated by a 1013 ms delay, was an identical match (DMS) or a category match (DMC). Because we observed similar results when we analyzed the DMS and DMC results separately (data not shown), we combined both tasks for all subsequent analysis. We refer to these two tasks as the motion task.

The roles of prefrontal and parietal cortices in encoding space and motion direction

As past studies have shown that both PFC and PPC are involved in the maintenance of task-relevant stimulus features in working memory, we compared how these two areas encode spatial location and motion direction during working memory. In 2 monkeys, we analyzed the activity of 343 PPC neurons (primarily from the LIP area, but potentially including areas 7a and the middle intraparietal area; see Materials and Methods) and 588 neurons from the dorsolateral PFC (arrays were centered over the principal sulcus). We only included neurons with sufficient number of trials from both tasks, had mean spike rates >1 Hz. Neurons were not prescreened for task-related responses or stimulus selectivity (see Materials and Methods).

We found that many PFC neurons were spatially and motion direction-selective ($p < 0.01$, one-way ANOVA) during stimulus presentation and during the delay. The percentages of neurons that were spatially or motion direction-selective during 333 ms windows covering the stimulus presentation, the middle, and late delay epochs are shown in Table 1. For example, the PFC neuron shown in Figure 2A preferentially responded to visual targets toward $\sim 0^\circ$ during the delay epoch of the memory saccade task, and for motion directions between $\sim 255^\circ$ and 315° during the delay epoch of the motion task. In contrast, many PPC neurons were spatially selective, but not motion direction-selective, during the middle delay (2.9% of PPC compared with 14.5% of PFC neurons were motion direction-selective, $p \sim 10^{-8}$, χ^2 test, $df = 1$; Table 1) and late delay (5.1% of PPC neurons compared with 14.5% of PFC neurons, $p \sim 10^{-5}$) epochs. For example, the PPC neuron shown in Figure 2B responded preferentially for spatial locations between $\sim 315^\circ$ and 0° during the delay epoch of the memory saccade task, but there is no obvious motion direction selectivity during the delay epoch of the motion task. This is despite clear motion direction selectivity during the stimulus presentation. Indeed, the percentage of PPC and PFC neurons selective for motion direction as measured during the stimulus presentation was approximately equal (20.3% of PPC neurons compared with 18.4% of PFC neurons, $p = 0.44$).

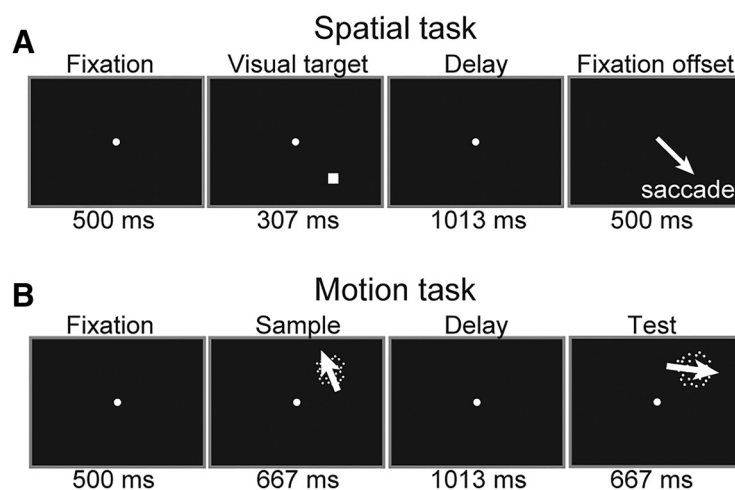


Figure 1. The spatial and motion tasks. **A**, The spatial task was a delayed memory saccade task, in which the monkeys had to maintain fixation on a central point for 500 ms, followed by a visual target presentation in one of eight locations for 307 ms, and finally by a delay period of 1013 ms. The central fixation point was then extinguished, and the monkey had 500 ms to saccade to the location of the remembered visual target to receive a reward. **B**, The motion task was either a DMS or a DMC task. The monkeys had to maintain fixation on a central point for 500 ms, followed by a 667 ms sample motion stimulus, followed by a 1013 ms delay, and then a 667 ms test motion stimulus. Monkeys released a manual lever to indicate whether the test stimulus was the same direction (DMS) or same category (DMC) as the previously presented sample. If the first test stimulus was a nonmatch, a second brief (267 ms) delay was presented followed by a test stimulus which matched the sample.

The percentages of PPC neurons that were selective for both spatial location and motion direction were 7.9%, 1.3%, and 1.9% for the stimulus, middle delay, and late delay epochs, respectively, and the percentages of PFC neurons that were selective for both was 7.1%, 5.3%, and 6.0% for the three epochs, respectively. These values were not significantly greater than the expected values if spatial location and motion direction selectivity were independent ($p > 0.05$ for all three epochs and both cortical areas, χ^2 test).

To further quantify these results across the neural population, we measured spatial location selectivity during the spatial task, and motion direction selectivity during the motion task, using linear SVM classifiers applied to pseudo-populations of PFC and PPC neurons (see Materials and Methods). We then calculated a decoding accuracy score, which measures how close the spatial location or motion direction predicted from the classifier was from the actual spatial location or motion direction. Values of 1 indicate perfect decoding, values of 0 indicate chance decoding, and values of -1 indicate that the predicted spatial or motion direction was 180° opposite was from the actual direction. For the spatial task (Fig. 3A), decoding accuracy for the PPC (green curve) and PFC (magenta curve) populations were significantly greater than chance ($p < 0.01$, bootstrap) from 60 and 80 ms after stimulus, respectively, until the end of the trial (Fig. 3A, top, horizontal green and magenta bars, times at which decoding accuracy was significantly greater than chance). Although the PPC decoding accuracy was significantly greater compared with PFC for a brief period during the stimulus presentation (Fig. 3A, bottom, horizontal green bars, $p < 0.01$, bootstrap), and the PFC decoding accuracy was significantly greater compared with PPC and intermittently during the delay epoch (Fig. 3A, bottom, horizontal magenta bars), spatial decoding accuracy was approximately comparable between PPC and PFC.

For the motion task, the PPC and PFC decoding accuracies followed markedly differing time courses. The PFC decoding accuracy was significantly greater than chance (Fig. 3B, top, horizontal magenta bars, $p < 0.01$, bootstrap) from 150 ms after

Table 1. Percentage of neurons selective for spatial location or motion direction^a

	Stimulus	Middle delay	Late delay
Spatial task			
PPC (<i>N</i> = 343)	30.0%	19.8%	23.6%
PFC (<i>N</i> = 588)	32.7%	24.8%	28.7%
χ^2	<i>p</i> = 0.41	<i>p</i> = 0.08	<i>p</i> = 0.09
Motion task			
PPC (<i>N</i> = 343)	20.3%	2.9%	5.1%
PFC (<i>N</i> = 588)	18.4%	14.5%	14.5%
χ^2	<i>p</i> = 0.44	<i>p</i> $\sim 10^{-8}$	<i>p</i> $\sim 10^{-5}$

^aPercentage of PPC and PFC neurons selective (one-way ANOVA, *p* < 0.01) for spatial location or motion direction during the stimulus epoch (80–412 ms after stimulus onset), the middle delay epoch (666–334 ms before the end of the delay), and the late delay epoch (the last 333 ms of the delay). *N* value indicates the number of valid PPC or PFC neurons. The χ^2 test compares the proportion of selective neurons between PPC and PFC for each task and epoch.

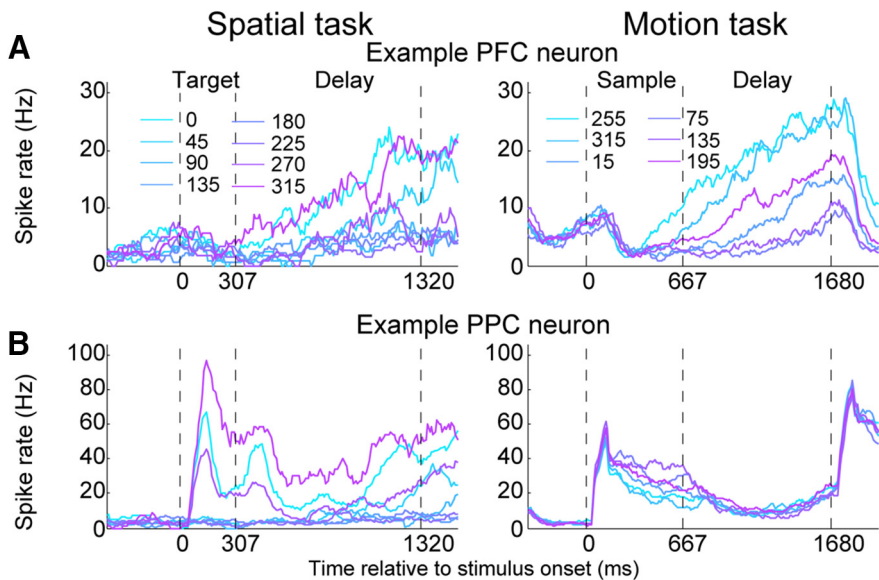


Figure 2. Example PFC and PPC neurons showing spatial and motion direction selectivity. **A**, Mean neural response of an example PFC neuron to the 8 spatial locations during the spatial task (left), and to the 6 motion directions during the motion task (right). Neural responses to the different spatial and motion directions are indicated by the different color traces. For the spatial task, the three vertical dashed lines, from left to right, indicate the visual target onset, the visual target offset, and the end of the delay period. For the motion task, the three vertical dashed lines, from left to right, indicate the sample stimulus onset, the sample offset, and the end of the delay period. **B**, Sample as **A**, except an example PPC neuron is shown.

stimulus onset until the end of the trial. In contrast, the PPC decoding accuracy was significantly greater than chance from 70 ms after stimulus onset until the early delay epoch (Fig. 3B, top, horizontal green bars), at which time its value became not significantly different from chance for the majority of the delay. The PPC decoding accuracy was initially greater than PFC during the early part of the sample epoch (Fig. 3B, bottom, horizontal green bars), before showing a large decrease during the delay epoch, during which time the PPC accuracy was significantly less than PFC (Fig. 3B, bottom, horizontal magenta bars). Thus, despite strongly encoding the motion stimulus during the sample period, PPC motion encoding significantly weakens during the delay.

We note that the weak delay period selectivity observed in Figure 3B is consistent with our past study showing that delay period selectivity in PPC during the DMC task emerges only after extensive categorization training (Sarma et al., 2016).

To confirm the results of this population decoding approach, we also measured the spatial and direction selectivity for each individual neuron by measuring the normalized percentage of variance of each neuron’s trial-by-trial spike rate that can be explained by the spatial location or motion direction (see Materials

and Methods). Greater values in the normalized PEV indicate that the firing rate of the neuron is increasingly selective for the spatial location or motion direction.

Consistent with the results above, spatial selectivity for both the PPC (Fig. 3C, left, green curve) and PFC (magenta curve) populations were significantly greater than chance (*p* < 0.01, bootstrap) from 60 and 80 ms after stimulus onset, respectively, until the end of the trial. PPC spatial selectivity was significantly greater compared with PFC during the early part of the stimulus presentation (Fig. 3C, bottom, horizontal green bars), and not significantly different at all other times.

Finally, motion selectivity time courses (Fig. 3D) were similar to the decoding accuracy time courses (Fig. 3B). Motion selectivity for the PPC and PFC neuronal populations were significantly greater than chance from 70 and 130 ms after stimulus onset, and PPC selectivity was significantly greater compared with PFC for the early part of the sample epoch (Fig. 3B, bottom, horizontal green bars), before dropping during the delay, during which time PPC selectivity was significantly less than PFC (Fig. 3B, bottom, horizontal magenta bars).

In summary, neurons in the PPC were only weakly motion direction-selective during the delay, despite strongly encoding the stimulus during the sample presentation. In contrast, motion direction selectivity in PFC and spatial location selectivity in PPC and PFC remained robust throughout the sample and delay epochs. This result begs the question of why the PPC fails to robustly represent motion direction information in working memory despite strong selectivity during the stimulus presentation, and strong delay-period spatial encoding during the spatial task.

Functional clustering and mnemonic encoding

We wanted to examine possible circuit mechanisms that could explain the difference in spatial location and motion direction mnemonic encoding in PPC and PFC. As stated in the Introduction, previous studies have suggested that stimulus-specific persistent activity is subserved by recurrent excitation among interconnected groups of neurons. Because (1) recurrent excitation within groups of interconnected neurons is strengthened when all neurons preferentially respond to the same stimulus, and (2) neurons are preferentially connected to other nearby neurons (Perin et al., 2011; Levy and Reyes, 2012), we hypothesize that persistent activity is facilitated by spatially clustered neurons that are similarly tuned. Measuring the presence of functional clustering was facilitated by our semichronic recording system, in which each electrode’s *x-y* coordinate position was fixed, but whose depth could be independently raised or lowered. This allowed us to accurately estimate the distance between neurons recorded in different recording sessions.

We measured the spatial and motion-direction functional clustering by calculating the weighted dot-product between the preferred spatial locations or motion directions of different neurons recorded on different days, which were located within 62.5–750 μ m. The preferred spatial locations and motion directions

were calculated by fitting the relationship between the stimulus and spiking activity with a linear model, and were weighted according to the proportion of variance explained by the linear model (see Materials and Methods). This tuning similarity measure could range between -1 and 1 , where a value of 1 indicates that nearby neurons prefer identical stimuli, 0 indicates no correlation between preferred stimuli, and -1 indicates that nearby neurons prefer opposite stimuli.

We should note that, in addition to local recurrent connections, both stimulus-specific bottom-up and top-down inputs can potentially contribute to functional clustering. Inferring the relative contribution of these sources is difficult, especially given that they may be context-dependent, and that their effects likely vary throughout the trial. Thus, for the analysis in Figure 4, we simply wish to measure the relationship between functional clustering and persistent activity agnostic of the source(s) of this functional clustering. In Figure 9, we will partly address the question of what sources contribute to functional clustering, as we examine how bottom-up spatial and motion input signals into PPC and PFC are organized.

For the spatial task (Fig. 4A), tuning similarity between nearby neurons was significantly greater than chance (top, bars, $p < 0.01$, bootstrap) both in PPC (green curve) and in PFC (magenta curve) throughout most of the stimulus and delay epochs. Thus, neurons that preferred similar spatial location were clustered in both areas, consistent with the hypothesized relationship between functional clustering and the presence of spatially selective persistent activity in both areas.

For the motion task (Fig. 4B), tuning similarity between nearby neurons in PFC (magenta curve) was significantly greater than chance throughout the sample and delay epochs. In contrast, tuning similarity in PPC (green curve) was above chance only intermittently throughout the sample and delay epochs, despite the fact that motion direction selectivity was actually stronger in PPC compared with PFC during the early part of the sample epoch (Fig. 3B,D). PFC tuning similarity was significantly greater than PPC during the latter part of the sample epoch and almost the entire delay (bottom, magenta bars). Thus, neurons that preferred similar motion directions were clustered in PFC, along with robust motion direction encoding through the delay period. In contrast, there existed significantly weaker clustering in PPC along with significantly weaker motion direction encoding during the delay.

Although these results show that mnemonic encoding and functional clustering are significantly different between cortical areas, it does not establish that neurons that are part of a cluster show stronger stimulus-specific persistent activity. Thus, we measured the correlation between a neuron's stimulus selectivity and the tuning similarity within sets of nearby neurons (see Materials and Methods). For example, if neurons A–D were all lo-

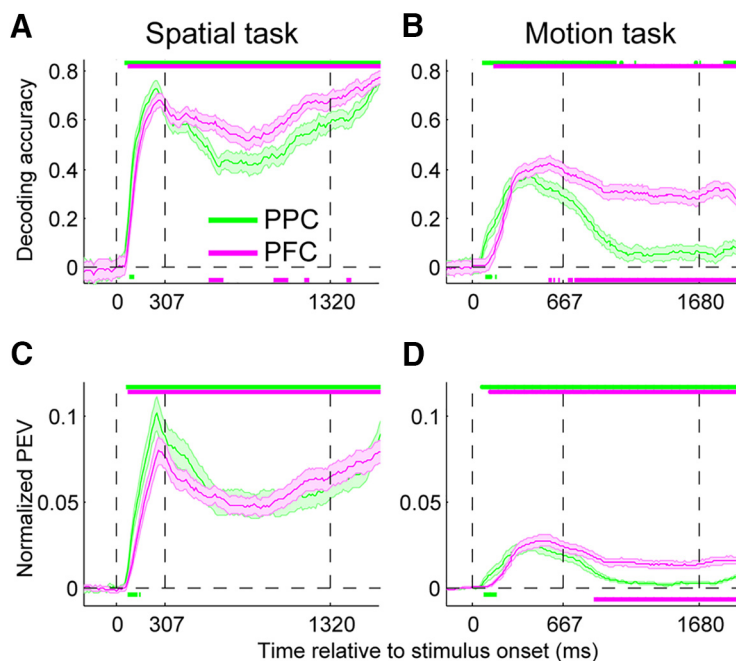


Figure 3. PFC and PPC population level spatial and motion direction selectivity. **A**, Decoding accuracy during the spatial task for the PPC (green curve) and PFC (magenta) neuronal populations was calculated by decoding the spatial location from the population spike rates using SVM linear classifiers. Top, Horizontal green bar represents times at which the mean PPC decoding accuracy was significantly greater than chance ($p < 0.01$, bootstrap). Horizontal magenta bar represents times at which the mean PFC decoding accuracy was significantly greater than chance ($p < 0.01$, bootstrap). Bottom, Horizontal green bar represents times at which the mean PPC decoding accuracy was significantly ($p < 0.01$, bootstrap) greater than the PFC decoding accuracy. Horizontal magenta bar represents times at which the mean PFC decoding accuracy was significantly ($p < 0.01$, bootstrap) greater than the PPC decoding accuracy. Error bars indicate SEM. **B**, Same as **A**, except for the motion task. **C**, Similar to **A**, except that spatial selectivity was calculated using the normalized PEV: the percentage of variance in each neuron's spike rate that is explained by the spatial or motion direction, normalized to eliminate positive bias arising from a finite number of samples. **D**, Same as **C**, except for the motion task.

cated within 62.5 – $750 \mu\text{m}$ from each other, we would first compare the stimulus selectivity of neuron A with the mean tuning similarity calculated between neurons B and C, B and D, and C and D. We would go on to compare the stimulus selectivity of neuron B with the tuning similarity between A, C, and D, etc. This correlation was calculated for all time points within a trial.

For the spatial task (Fig. 4C), tuning similarity was significantly correlated with stimulus selectivity for parts of the stimulus presentation and most of the delay epoch in PPC (Fig. 4C, top, horizontal green bars, $p < 0.01$, bootstrap) and PFC (Fig. 4C, top, horizontal magenta bars). Thus, for PPC and PFC, neurons that mnemonically encoded the spatial location tended to be part of a cluster in which neurons are tuned to similar spatial locations. We do note that the correlation for both PPC and PFC decreased immediately after stimulus offset, which is potentially because visually selective neurons become less active while those more involved in preparing the oculomotor response or maintaining the spatial location in working memory become more active (Markowitz et al., 2015). Alternatively, from a dynamical systems standpoint, the stimulus offset removes an attractor manifold and causes neural trajectories in state space to drift toward a different attractor (Chaisangmongkon et al., 2017).

For the motion task, we also observed a significant correlation between tuning similarity and motion direction selectivity in PFC (Fig. 4D, top, horizontal magenta bars), through latter part of the sample and the entire delay epoch. In PPC, the correlation between tuning similarity and motion direction selectivity in PPC was seldom significantly >0 at any point during the sample or

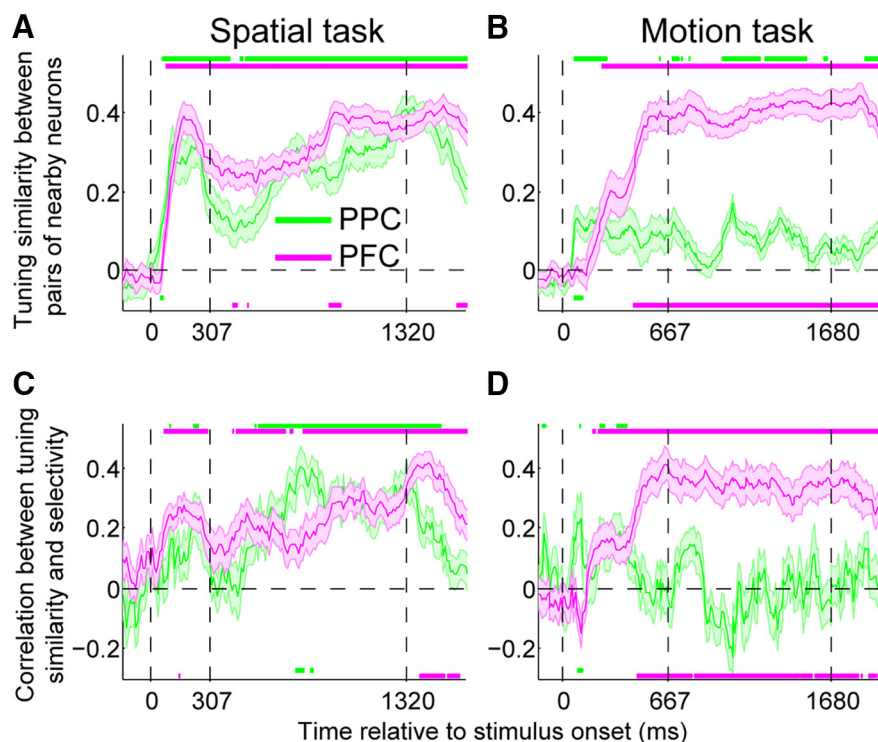


Figure 4. Tuning similarity and persistent activity in PPC and PFC. **A**, Spatial tuning similarity for the PPC (green curve) and the PFC (magenta) populations was calculated as the weighted dot-product between the preferred spatial location or motion directions of nearby neurons recorded during different sessions. As with Figure 2, vertical dashed lines indicate, from left to right, the stimulus presentation, the stimulus offset, and the end of the delay period. Error bars indicate SEM. **B**, Same as **A**, except for the motion task. **C**, The time course of the correlation between the spatial tuning similarity and the spatial stimulus selectivity (i.e., normalized PEV). **D**, Same as **C**, except showing the correlation between the motion direction tuning similarity with motion direction selectivity.

delay. Furthermore, the correlation in PFC was significantly greater than the correlation in PPC for part of the sample and the entire delay epoch (Fig. 4D, bottom, horizontal magenta bars). This was not unexpected given PPC's weak motion direction tuning similarity.

To ensure that these results were consistent across monkeys, we repeated our analysis for each individual subject (Fig. 5). For Monkey Q, PPC motion decoding accuracy (Fig. 5A), tuning similarity (Fig. 5B), and correlation between tuning similarity and selectivity (Fig. 5C) were all not significantly greater than chance for almost all of the delay epoch. In Figure 5D, E, we show the scatter plot of all individual clusters of nearby neurons used to calculate the correlations in Figure 4C. Averaging the tuning similarity (x -axis) and stimulus selectivity (y -axis) across the last 500 ms of the delay epoch, we find that tuning similarity and stimulus selectivity are positively correlated in all cases, except in PPC for the motion task, also consistent with our above results.

The results of Monkey W are mostly similar, except that PPC motion decoding accuracy was somewhat stronger, and was intermittently greater than chance during the delay epoch (Fig. 5F). Consistent with this, PPC motion direction tuning similarity was also significantly greater than chance during most of the delay epoch, although it was still weaker than PFC motion direction tuning similarity during parts of the delay (Fig. 5G). And similar to Monkey Q, the correlation between tuning similarity and stimulus selectivity averaged across the last 500 ms of the delay epoch is significantly >0 in PFC for the spatial and motion tasks (Fig. 5J), and in PPC during the spatial task, but not the motion task (Fig. 5I).

Tuning similarity of as a function of distance

We wondered whether the lack of clustering of PPC neurons based on motion direction preference was because our distance range (i.e., we compared all pairs of neurons spaced between 62.5 and 750 μm apart) was too large. It is possible that motion direction clustering in PPC operates on a much finer scale compared with PFC, or compared with clustering based on preferred spatial location. Thus, in Figure 6, we repeated our analysis used in Figure 4A, B, except that we used three distance ranges: we calculated tuning similarity for pairs of neurons that were spaced between 62 and 188 μm apart (0.5–1.5 turns, green curves), between 188 and 469 μm (1.75–3.75 turns, magenta curves), and between 469 and 750 μm (4–6 turns, cyan curves).

For the spatial task in PPC (Fig. 6A), tuning similarity for the nearest distance range (green curve) was significantly greater throughout the stimulus presentation and delay epoch. While the tuning similarities for the two furthest distance ranges (magenta and cyan curves) were not as strong, they were still significantly greater than chance throughout large portions of the stimulus presentation and delay epochs.

For the spatial task in PFC (Fig. 6B), tuning similarity for the two nearest distance ranges (green and magenta curves) were significantly greater throughout the stimulus presentation and delay epoch. While the tuning similarity for the furthest distance range (cyan curve) was not as strong, it was still significantly greater than chance throughout large portions of the stimulus presentation and delay epochs.

For the motion task in PPC (Fig. 6C), tuning similarity for the nearest distance range (green curve) was significantly greater during the stimulus presentation and early delay. In contrast, tuning similarity for the two furthest distance range (magenta and cyan curve) were not significantly >0 for almost the entire trial. This was the case for both monkeys individually (data not shown).

For the motion task in PFC (Fig. 6D), the strength and latency of tuning similarity appeared to vary according to the distance range: tuning similarity was stronger, and developed with shorter latency, for nearer distance ranges.

These results suggest several hypotheses. First, and unsurprising, is that there is a tendency for tuning similarity to decrease as the distance between neurons increases. Second, for those cases in which stimulus information is robustly encoded during working memory (spatial task for PPC and PFC, motion task for PFC), tuning similarity was significantly greater than chance during the delay, even for the greatest distance between pairs of neurons that we examined. This suggests that persistent activity is associated with functional clusters existing on scales of $>500 \mu\text{m}$. In contrast, PPC neurons were clustered based on their preferred motion direction, but these clusters were on scales of $<200 \mu\text{m}$,

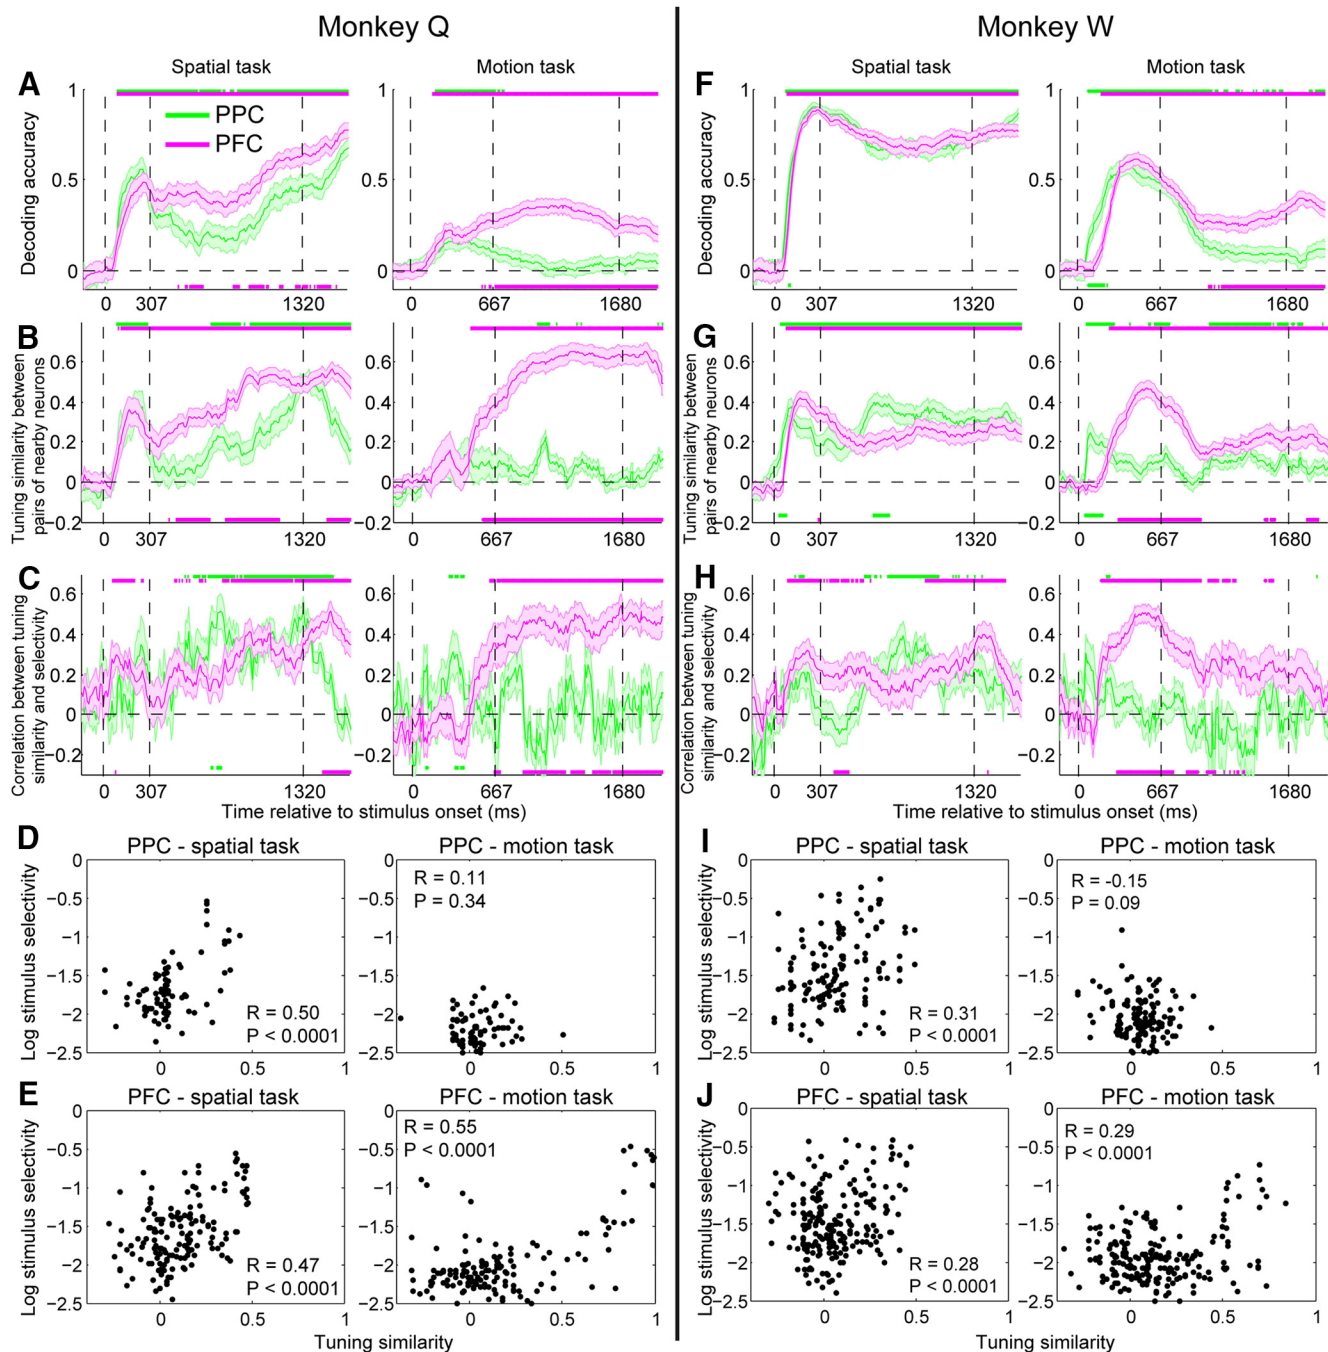


Figure 5. Tuning similarity and persistent activity in PPC and PFC for individual monkeys. **A**, Similar to Figure 3A, the decoding accuracy of PPC (green curve) and PFC (magenta curve) is shown for Monkey Q during the spatial (left) and motion (right) tasks. **B**, Similar to Figure 4A, B, the mean tuning similarity is shown for Monkey Q. **C**, Similar to Figure 4C, D, the time course of the correlation between the spatial tuning similarity and the spatial stimulus selectivity is shown for Monkey Q. **D**, Scatter plot showing the individual clusters of nearby neurons that are used to compute the correlation in **C** for Monkey Q. The tuning similarity (x-axis) and the logarithm of the stimulus selectivity (y-axis) were averaged over the last 500 ms of the delay epoch for PPC for the spatial (left) and motion (right) tasks. **E**, Same as **D**, except for the motion task. **F–J**, Same as **A–E**, except for Monkey W.

consistent with the view that functional clusters must be large enough to support working memory.

Recording locations and depths

We wondered whether the difference between PPC and PFC described in Figures 3–5 was because of any biases in our recording locations or depths. For example, a previous study found that PFC neurons spatially selective during working memory are primarily located posterior and lateral to the principal sulcus and at more shallow recording depths (Markowitz et al., 2015). Al-

though it is in theory possible that lack of clustering of preferred motion directions in PPC was because of biases in our neuronal recording approach, several lines of evidence argue against this possibility. First, our PPC recordings revealed strong organization of spatial selectivity, implying that the lack of clustering of preferred motion directions was not because we did not record from areas that support persistent delay-period activity in general. Second, we recorded from neurons over a wide span of PPC and PFC locations from both animals (Fig. 7A–D), with 31 of 64 PPC recording locations across 2 monkeys contributing five or

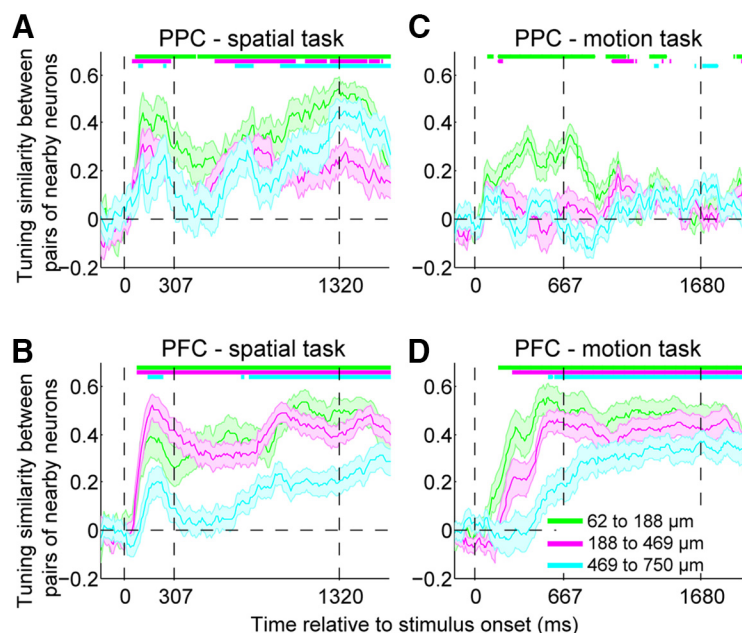


Figure 6. Tuning similarity as a function of distance. Tuning similarity was calculated in the same manner as done for Figure 4A, B, except that similarity was calculated between all pairs of neurons spaced between 62 and 188 μm apart (green curve), between 188 and 469 μm apart (magenta curve), and between 469 and 750 μm apart (cyan curve). **A**, The mean tuning similarity for the three distance ranges for PPC during the spatial task. Top, Horizontal bars represent times at which the tuning similarity was significantly greater than chance ($p < 0.01$, bootstrap). **B**, Same as **A**, except for PFC during the spatial task. **C**, Same as **A**, except for PPC during the motion task. **D**, Same as **A**, except for PFC during the motion task.

more neurons toward our analysis. Third, recent studies have suggested that neurons encoding stimulus information during working memory might be preferentially located in superficial layers (Wang et al., 2013; Markowitz et al., 2015). While unfortunately our recording approach did not have the precision needed to confidently determine the cortical layer of recorded neurons, the distribution of estimated recording depths (see Materials and Methods; Fig. 7E) suggests that we broadly sampled across cortical layers, and if anything, PPC recordings were more biased toward superficial layers than PFC.

Last, we wondered whether motion direction tuning similarity in PPC was weaker because pairs of nearby PPC neurons were spaced further apart compared with pairs of PFC neurons. However, the opposite was true: the mean distance between pairs of PPC neurons used for the analysis in Figure 4A, B was 360 μm compared with 392 μm in PFC ($p = 0.001$, two-tailed, two-sample t test, $df = 1835$, Fig. 7F). In summary, we not believe that the differences between PPC and PFC can be explained by biases in our neuronal recording approach.

Mnemonic encoding in a trained recurrent neural network model

We wanted to more formally test our hypothesis that persistent activity is facilitated when groups of interconnected neurons are similarly tuned. To do so, we wanted to examine persistent activity in neural network models in which we could control the connectivity between different groups of neurons. Thus, we trained recurrent neural networks using the Pycog framework (Song et al., 2016) to perform a variant of the DMS task described above (see Materials and Methods). The model (Fig. 8A) consisted of 72 motion direction-selective neurons that randomly projected onto a network of 180 recurrently connected neurons. Neurons were organized into 6 “columns” consisting of 24 excitatory and 6 inhibitory neurons; 50% of pairs of neurons within a column

were connected, whereas 5% of pairs of neurons from different columns were connected. Neurons from this recurrently connected network randomly projected onto two decision neurons: a “match” neuron, which was trained to signal when the sample and test directions matched, and a “nonmatch” neuron, which was trained to signal when the sample and test directions did not match. All connection weights were randomly initialized, so the preferred neurons’ preferred motion directions within the recurrent network were initially completely random. To ensure that our results generalized across networks sharing this topology, we trained 20 different neural networks initiated with different random weight matrices, and included all the resultant data into the following analysis.

The network was trained so that the “match” neural response increased from 0.2 to 1.0 (values arbitrarily set) for a matching test stimulus, and that the “nonmatch” neural response increased from 0.2 to 1.0 for a nonmatching test stimulus. We successfully trained the network to achieve these targets, as the mean match neural response (green curve) selectively increased from ~ 0.2 to 1.0 for a matching test stimulus (Fig. 8B), and the mean nonmatch neural response (magenta curve) selectively increased from ~ 0.2 to 1.0 for a nonmatching test stimulus (Fig. 8C).

Given that the model was successfully trained to perform the delayed matching task, we wanted to know (1) whether neurons within a column developed similar direction tuning, and (2) whether columns with more similar tuning among its neurons also had greater persistent activity.

To answer these questions, we analyzed the network in a similar fashion to how we analyzed our experimental data (Figs. 3–6). After training, model neurons within the recurrent network selectively encoded the motion direction during both the sample and delay epochs (Fig. 8D). Concurrent with this selective encoding of the motion direction, tuning within columns self-organized: direction tuning was significantly more similar between neurons within the same column (Fig. 8E, blue curve) compared with pairs of neurons from different columns (Fig. 8E, red curve). Finally, columns with more similar tuning among its neurons more selectively encoded the motion stimulus; Figure 8F is a scatter plot comparing the mean motion direction selectivity of each column measured at the end of the delay (x -axis) versus the tuning similarity within each column, also measured at the end of the delay (y -axis). At this time point, the correlation between both features is $r = 0.50$ ($p < 0.001$, $df = 118$). This correlation between motion direction selectivity and tuning similarity within each column develops during the sample stimulus and is maintained throughout the entire delay (Fig. 8G).

Thus, in recurrent neural networks embedded with a columnar topology, training on the DMS task changes the connection weights so that neurons within the same column develop similar motion direction tuning. Furthermore, the level of tuning similarity within a column correlates with the level of stimulus-

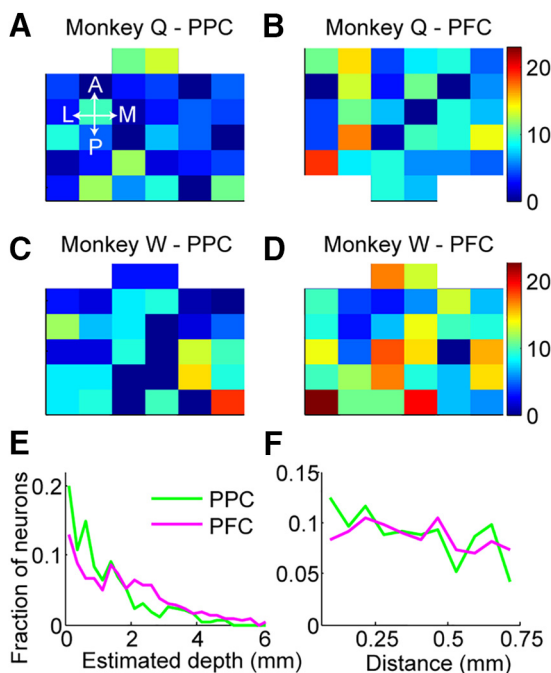


Figure 7. Recording locations, depths, and distances between neuron pairs. **A**, Heat map showing the number of neurons included in our analysis at each recording location on the array placed over PPC on Monkey Q. The array consisted of 32 electrodes, with 5 rows of 6 electrodes, with a final row of 2 electrodes at the anterior edge. **B**, Same as **A**, except for the array placed over PFC on Monkey Q. **C**, Same as **A**, except for the array placed over PPC on Monkey W. **D**, Same as **A**, except for the array placed over PFC on Monkey W. **E**, Histograms showing the distributions of estimated cortical depths for all PPC (green curve) and PFC (magenta curve) neurons included in our analysis. The cortical depth was estimated by finding the shallowest depth of all neurons recorded from each channel, and setting that depth equal to zero. **F**, Histograms showing the distributions of distances between neuron pairs for PPC (green curve) and PFC (magenta curve). As per our analysis, only neuron pairs separated by distances within a range of 62.5–750 μm are included.

selective persistent activity, confirming our intuition that cortical organization supports working memory.

Stimulus selectivity in the LFP

In Figures 4–6, we established that mnemonic encoding is associated with the presence of functional clustering. Is this functional clustering (partially) the result of organized spatial and motion input signals arriving into these areas, or does clustering occur after the input signals arrive? Although we cannot directly measure the inputs into these areas, we can infer their organization by examining responses in the LFP. The LFP is thought to mainly reflect the sum of synaptic potentials within several hundred microns of the electrode tip (Katzner et al., 2009; Kajikawa and Schroeder, 2011). If the distribution of synaptic inputs selective for various stimuli or features was randomly distributed within a cortical area, the mean synaptic activity within a local volume would be equal across stimuli, leading to no selectivity in the LFP response. However, if synaptic inputs selective for each stimulus tend to be in close spatial proximity (i.e., functionally clustered), the summed synaptic activity within a local volume would vary across stimuli, leading to selectivity in the LFP response.

In Figure 9, we show the evoked potentials (i.e., the mean LFP relative to the stimulus presentation, causally filtered with a 10 ms boxcar) for the spatial (Fig. 9A) and motion (Fig. 9B) tasks for the same example electrode. For this electrode, the evoked potential follows a similar time course in both tasks, with weak depolarization (negative deflection) ~ 60 ms after stimulus presentation,

followed by hyperpolarization at ~ 90 ms, followed by stronger depolarization peaking at ~ 140 ms. The evoked potentials appear selective for spatial location, with stronger hyperpolarization (positive deflections) for spatial locations at 0 and 45 degrees, followed by stronger depolarizations for stimuli at 0, 45, and 90 degrees. In contrast, motion direction selectivity appears much weaker, with slightly stronger hyperpolarization for the motion direction at 255 degrees, followed by a slightly stronger depolarization for motion directions at 135 and 195 degrees.

To quantify the selectivity across the population of LFP recordings, we used linear SVMs to decode spatial locations and motion directions (Fig. 9C, similar to our approach used in Fig. 3A,B; see Materials and Methods). Consistent with the example electrode shown in Figure 9A, spatial selectivity in PPC (blue curve) becomes significantly greater ($p < 0.01$, bootstrap, indicated by the blue horizontal bars) than zero at 60 ms after stimulus presentation, and at 70 ms in PFC (red curve). Motion direction selectivity developed later, reaching significance in PPC (green curve) at 250 and 280 ms in PFC (black curve). Furthermore, maximum spatial decoding accuracy was ~ 10 times greater than maximum motion direction decoding accuracy in both PPC and PFC and was significantly greater ($p < 10^{-6}$, bootstrap) from 70 to 230 ms in PPC, and from 80 to 230 ms in PFC. These results suggest that spatial location signals arriving in both PPC and PFC are already functionally clustered, whereas motion direction signals arriving in PPC and PFC are significantly less organized.

Discussion

Although stimulus-specific persistent activity in the frontoparietal cortex is commonly observed during tasks requiring working memory (Funahashi et al., 1989; Colby et al., 1996; Chafee and Goldman-Rakic, 1998; Rainer et al., 1998; Romo et al., 1999; Zaksas and Pasternak, 2006), it is not understood why it is present for some, but not all, visual features. In this study, we seek to understand the neuronal conditions that can potentially allow for persistent activity in certain contexts and inhibit it in others.

To address this question, we showed that although PPC and PFC robustly encode spatial location during the working memory delay period of the task, only the PFC shows robust working memory encoding of motion direction. We found that this difference in mnemonic encoding could be partially explained by differences in functional clustering of spatial and motion selectivity: pairs of neurons in the PPC or PFC neurons within ~ 700 μm of each other preferred similar spatial locations, and pairs of PFC neurons within ~ 700 μm of each other preferred similar motion directions. In contrast, the preferred motion directions of PPC neuron pairs separated by more than ~ 200 μm were not correlated. We conclude that functional clustering facilitates persistent activity and can potentially explain why persistent activity is present for some, but not all, visual features.

Emergence of category selectivity during working memory in PPC during category learning

In a previous study, we showed that PPC mnemonically encodes the category membership of a motion direction stimulus after extensive categorization training (Sarma et al., 2016). In this study, we recorded across ~ 30 sessions while the monkeys were trained to perform the DMC categorization task. This proved to be insufficient time for the monkeys to reach a high level of categorization performance using this training approach, which likely explains why robust mnemonic motion-category encoding failed to develop in PPC, similar to what we observed in a recent

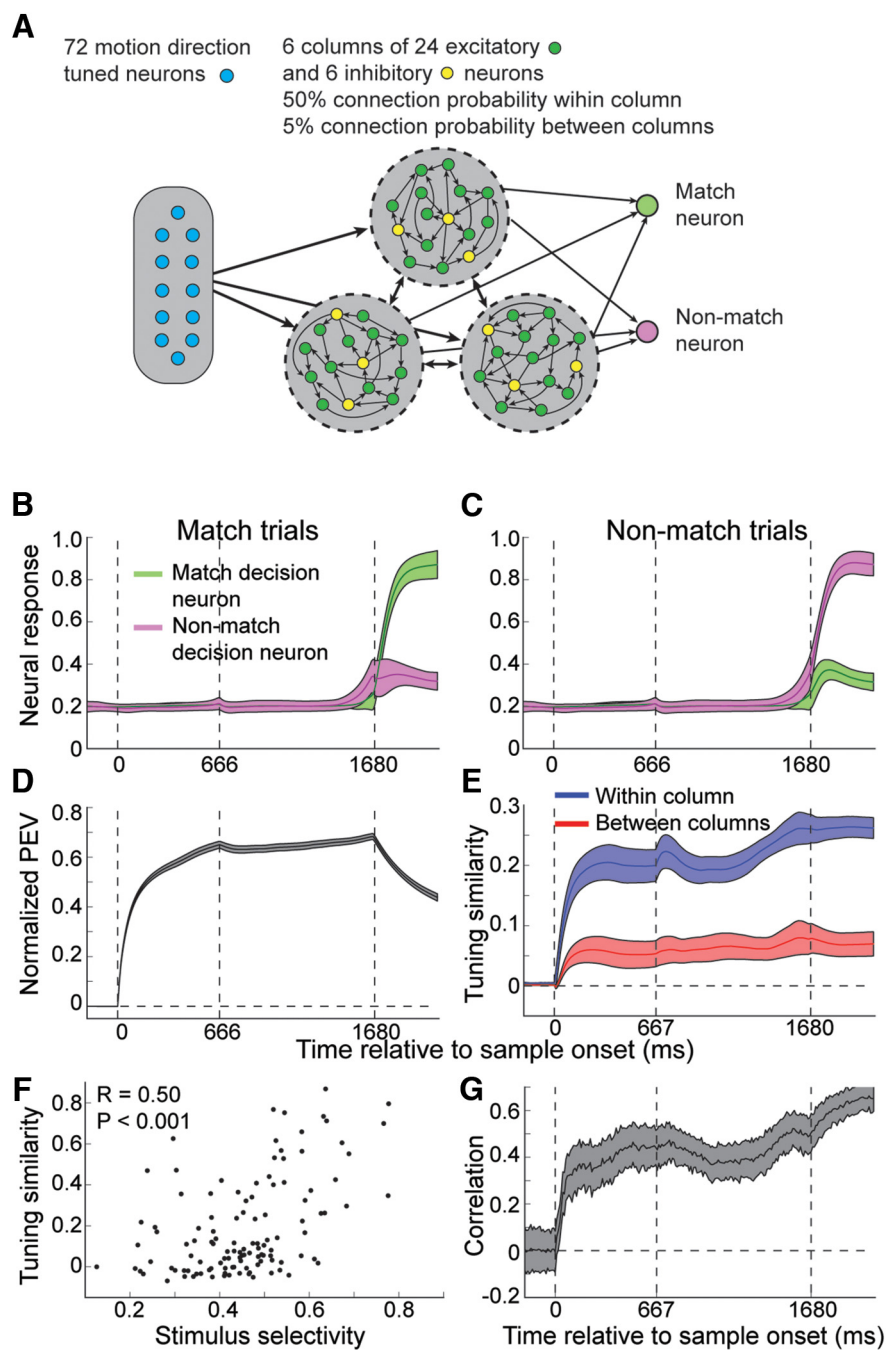


Figure 8. Direction encoding and connectivity in trained recurrent neural network models. **A**, Architecture of the neural network trained to perform the DMS task. Input consisted of 72 excitatory motion direction-selective neurons whose preferred directions uniformly spanned 360 degrees. Each input neuron had a 50% connection probability with each neuron in the downstream recurrent network. The recurrent network consisted of 6 columns of 24 excitatory and 6 inhibitory neurons (180 neurons total). Connection probability within a column was 50%, and the connection probability between neurons from different columns was 5%. These recurrent connected neurons projected onto two downstream output neurons that were trained to selectively respond to a matching test stimulus or to a nonmatching test stimulus. The connection probability between each neuron in the recurrent network and each downstream neuron was 50%. **B**, The mean response of the match (green curve) and nonmatch (magenta curve) output neurons for match trials. Mean responses were averaged across all 20 trained networks. Error bars indicate 1 SD. **C**, Same as **B**, except for nonmatch trials. **D**, The normalized PEV averaged across all 3600 neurons in the recurrent network (20 networks each containing 180 recurrently connected neurons). Error bars indicate SEM. **E**, The mean motion direction tuning similarity across all pairs of neurons within the same column (blue curve) and all pairs of neurons from different columns (red curve). **F**, Scatter plot comparing the mean motion direction selectivity for each column measured using the normalized PEV (x-axis) versus the mean tuning similarity (y-axis) for all pairs of neurons within the same column. Both values were calculated from the last time point in the delay epoch. Each dot represents one of the 120 columns (20 networks each containing 6 columns). **G**, The time course of the correlation between the mean motion direction selectivity (normalized PEV) for each column versus the mean tuning similarity of neurons within the same column.

study (Sarma et al., 2016). Thus, future studies will be required to address the role of functional clustering in the development of training-dependent persistent activity.

Innate or context-dependent feature maps

Is the functional clustering that we observe context-dependent, or is it innately hard-coded through long-term synaptic changes in the frontoparietal network? In Figure 9, we analyzed the spatial and motion evoked potentials to infer that the spatial inputs into PPC and PFC are clustered according to their preferred spatial locations, but that motion inputs did not show clustering according to their preferred directions. However, we should point out several possible caveats with this conclusion. First, this inference rests upon the assumption that the evoked potentials reflect the mean synaptic input within a small local volume (Katzner et al., 2009; Kajikawa and Schroeder, 2011). If motion-selective inputs into PPC and PFC arrive later than spatial inputs, they might be masked by the large deflections in the evoked potential generated by the stimulus onset. Third, there is no direct way to compare motion and spatial selectivity, and measuring spatial selectivity using eight spatial locations, while measuring motion selectivity with only six directions, might bias our results.

Despite these caveats, human imaging studies have also suggested that spatial clustering, in the form of retinotopic maps, is likely innate in PPC, as several subregions of the parietal cortex are retinotopically organized even for passively viewed stimuli (Swisher et al., 2007). Although few studies have examined whether spatial maps exist in PFC under passive conditions, two nearby frontal areas, the frontal and supplementary eye fields, are both topologically connected to areas in visual cortex (Schall et al., 1993, 1995), implying that, at a minimum, spatial maps are innately present within certain segments of the frontal cortex. These results are consistent with the known role that both areas (particularly the PPC) play in spatial processing (Colby and Goldberg, 1999; Constantinidis, 2006; Bisley and Goldberg, 2010).

Although innate spatial organization might be present in PPC and PFC, it is less certain that both areas are organized according to motion direction. Motion direction tuning similarity in PFC (Fig. 4B) requires >500 ms to approach its maximum value, and motion direction LFP se-

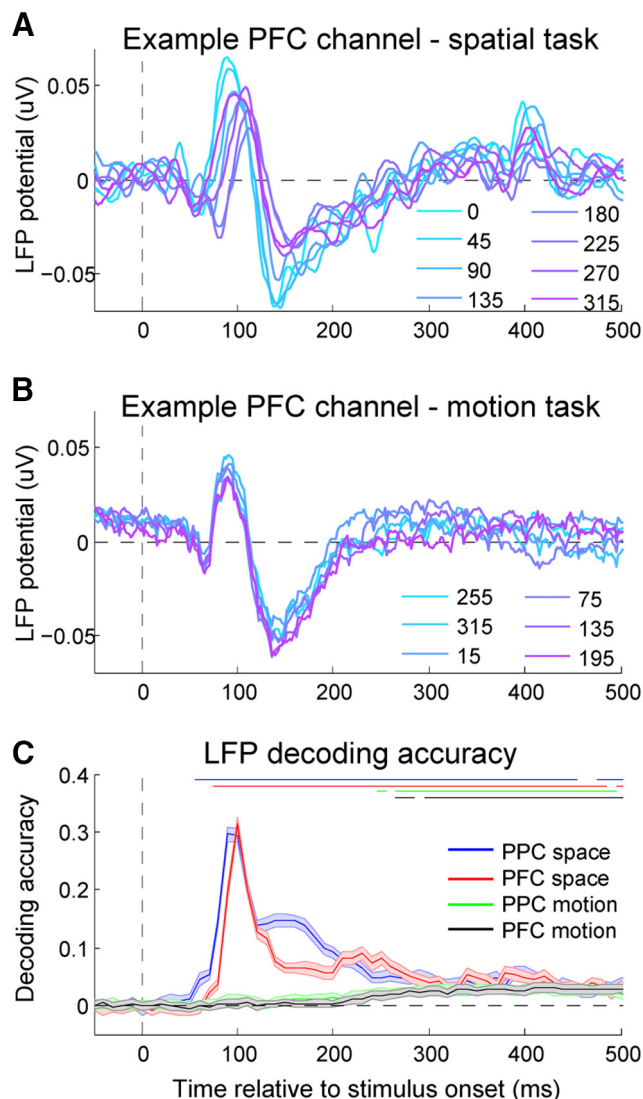


Figure 9. LFP spatial and motion direction selectivity. **A**, The mean LFP response, filtered with a 10 ms boxcar, to the eight target locations during the spatial task from an example electrode in PFC. Vertical dashed line indicates the time of visual target onset. **B**, The mean LFP response to the six motion directions during the motion task from the same example electrode. Vertical dashed line indicates the time of sample onset. **C**, LFP decoding accuracy during the spatial task for PPC (blue curve) and PFC (red), and during the motion task for PPC (green) and PFC (black) was calculated by decoding the spatial location from the LFP from all electrodes during each session using SVM linear classifiers. Top, Horizontal bars represent times at which the decoding accuracy was significantly ($p < 0.01$, bootstrap) greater than chance. Vertical dashed line indicates the time of visual target (spatial task) or sample (motion task) onset. Error bars indicate the SEM.

lectivity (Fig. 9C) is weaker than spatial selectivity. This stands in contrast to the rapid (<100 ms) development of spatial location tuning similarity in PPC and PFC (Figs. 4A, 9C), suggesting that more dynamic processes might be involved in forming clusters based on preferred motion directions. Unfortunately, because we did not map spatial or motion selectivity outside the context of the spatial and motion tasks in this experiment, it is difficult to provide a definitive answer to this question. However, assuming that clustering for motion direction is not innate, then understanding how functional clustering is regulated in a context-dependent manner is an important next question.

Regulation of persistent activity through spike-LFP interactions

The possible dynamic regulation of functional clustering in support of working memory would be in line with several other studies that have examined how working memory is subserved by short-term synaptic plasticity (Mongillo et al., 2008; Stokes, 2015). The proposal is that persistent activity results from the spiking of neural ensembles that are transiently assembled through short-term synaptic modifications (Harris, 2005; Fujisawa et al., 2008; Szatmáry and Izhikevich, 2010). In one interpretation, this implies that the contents of working memory are not stored in the patterns of neural activity per se, but rather within these short-term synaptic changes (Mongillo et al., 2008; Stokes, 2015). Although measuring these possible synaptic changes is challenging with extracellular recordings, these short-term synaptic changes might manifest themselves through their interactions with oscillations in the LFP. Recent studies have proposed that these short-term synaptic changes increase neurons' phase-locking to LFP oscillations, most likely within the gamma (40–100 Hz) frequency band (Lundqvist et al., 2016). Unfortunately, our recording approach is ill suited to examine this possibility: to avoid spiking contamination of the LFP phase measurement, it is advisable to measure how spikes on one electrode interact with the LFP on a different electrode. In our system, electrodes are spaced 1.5 mm apart, whereas gamma oscillations are thought to be a more local phenomenon (Buzsáki and Wang, 2012). However, working memory is also likely supported by more long-range connections, which may be facilitated by LFP oscillations in the alpha (Jensen et al., 2002; Foster et al., 2016) and beta (Salazar et al., 2012) frequency bands. How mnemonic encoding of spatial location and motion direction are subserved by spike-LFP interactions in these lower frequency bands will be a question for future research.

The role of persistent activity

Although this study examines the mechanisms that facilitate persistent activity, it does not address the role of persistent activity in working memory. Historically, persistent activity in the fronto-parietal cortex has been viewed as the substrate in which remembered information is stored (Funahashi et al., 1989; Colby et al., 1996; Chafee and Goldman-Rakic, 1998; Rainer et al., 1998; Romo et al., 1999; Ester et al., 2015). However, more recent work has suggested that this activity reflects the coordination with sensory and motor areas to form an appropriate behavioral response (Harrison and Tong, 2009; Lara and Wallis, 2014; Sreenivasan et al., 2014; Pasternak et al., 2015), or alternatively, that both storage and behavioral response-related modes exist in distinct zones (Markowitz et al., 2015). The existence of a response-related mode would suggest that persistent activity during the delayed memory saccade task might be more related to the deployment of spatial attention toward the target, or the preparation of the upcoming oculomotor response. In the motion task, persistent activity might reflect network changes that will allow for the proper behavioral response to the upcoming test stimulus (i.e., a prospective code) (Rainer et al., 1999; Stokes et al., 2013). This raises the possibility that the association between functional clustering and persistent activity generalizes beyond the context of working memory: various cognitive processes, such as spatial attention and motor planning, in which changes in neural activity must persist from the stimulus presentation until the motor response, might also require the formation task-related cortical organization (Ikkai and Curtis, 2011).

It might also explain a recent study that showed that stimulus-selective persistent activity was absent from PFC during a fine-

color change detection task (Lara and Wallis, 2014), and the results of other studies showing that PFC encodes abstract or category-like variables more accurately than variables relating to image similarity (Freedman et al., 2003; Meyers et al., 2008). Cortical circuits in PFC might not be organized in a manner that allows for clustering based on the fine gradation of a visual feature, and must instead coordinate with visual cortex to encode precise visual details.

Recurrent neural network models

Until recently, attempts to train recurrent neural networks often failed because, when trying to minimize the loss function, the gradient had a tendency to either explode or approach zero (Pascanu et al., 2012). Recent advances in training algorithms have circumvented this problem, allowing researchers to train recurrent networks to perform various computations (Mante et al., 2013; Song et al., 2016; Chaisangmongkon et al., 2017). In Figure 8, we trained recurrent neural network models, in which neurons within the same “column” were preferentially connected, to solve a delayed matching task. Consistent with the experimental data, neurons within the same column were more similarly tuned compared with pairs of neurons in different columns, and the level of tuning similarity in each column was correlated with the stimulus selectivity of those neurons throughout the delay epoch.

Although these models are gross simplification of actual neural circuits, there has been growing interest in using artificial neural networks to confirm hypotheses from neural recordings or to gain insight into how the brain performs certain computations (Mante et al., 2013; Song et al., 2016; Chaisangmongkon et al., 2017). A major advantage of these models is that one can analyze the synchronous activity of the entire model population, a feat that is not possible with real experimental data (Mante et al., 2013). This allows one to fully reverse-engineer how the artificial circuits implement various computations.

In conclusion, we have proposed that mnemonic encoding of visual stimuli is facilitated by functional clustering, in which nearby neurons preferentially respond to similar stimuli. The inability of PPC or PFC to form appropriate functional clusters in certain contexts and/or for certain visual stimuli helps explain why persistent activity exists in these areas in some, but not all, cases. Future research will examine whether functional clustering is dynamically regulated, and the mechanisms that underlie its emergence.

References

- Asaad WF, Santhanam N, McClellan S, Freedman DJ (2013) High-performance execution of psychophysical tasks with complex visual stimuli in MATLAB. *J Neurophysiol* 109:249–260. [CrossRef Medline](#)
- Bisley JW, Goldberg ME (2010) Attention, intention, and priority in the parietal lobe. *Annu Rev Neurosci* 33:1–21. [CrossRef Medline](#)
- Buschman TJ, Siegel M, Roy JE, Miller EK (2011) Neural substrates of cognitive capacity limitations. *Proc Natl Acad Sci U S A* 108:11252–11255. [CrossRef Medline](#)
- Buzsáki G, Wang XJ (2012) Mechanisms of gamma oscillations. *Annu Rev Neurosci* 35:203–225. [CrossRef Medline](#)
- Buzsáki G, Anastassiou CA, Koch C (2012) The origin of extracellular fields and currents: EEG, ECoG, LFP and spikes. *Nat Rev Neurosci* 13:407–420. [CrossRef Medline](#)
- Chafee MV, Goldman-Rakic PS (1998) Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *J Neurophysiol* 79:2919–2940. [Medline](#)
- Chaisangmongkon W, Swaminathan SK, Freedman DJ, Wang XJ (2017) Computing by robust transience: how the fronto-parietal network performs sequential, category-based decisions. *Neuron* 93:1504–1517.e4. [CrossRef Medline](#)
- Colby CL, Goldberg ME (1999) Space and attention in parietal cortex. *Annu Rev Neurosci* 22:319–349. [CrossRef Medline](#)
- Colby CL, Duhamel JR, Goldberg ME (1996) Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *J Neurophysiol* 76:2841–2852. [Medline](#)
- Constantinidis C (2006) Posterior parietal mechanisms of visual attention. *Rev Neurosci* 17:415–427. [Medline](#)
- Ester EF, Sprague TC, Serences JT (2015) Parietal and frontal cortex encode stimulus-specific mnemonic representations during visual working memory. *Neuron* 87:893–905. [CrossRef Medline](#)
- Foster JJ, Sutterer DW, Serences JT, Vogel EK, Awh E (2016) The topography of alpha-band activity tracks the content of spatial working memory. *J Neurophysiol* 115:168–177. [CrossRef Medline](#)
- Freedman DJ, Assad JA (2006) Experience-dependent representation of visual categories in parietal cortex. *Nature* 443:85–88. [CrossRef Medline](#)
- Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2003) A comparison of primate prefrontal and inferior temporal cortices during visual categorization. *J Neurosci* 23:5235–5246. [Medline](#)
- Fujisawa S, Amarasingham A, Harrison MT, Buzsáki G (2008) Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat Neurosci* 11:823–833. [CrossRef Medline](#)
- Funahashi S, Bruce CJ, Goldman-Rakic PS (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol* 61:331–349. [Medline](#)
- Harris KD (2005) Neural signatures of cell assembly organization. *Nat Rev Neurosci* 6:399–407. [CrossRef Medline](#)
- Harrison SA, Tong F (2009) Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458:632–635. [CrossRef Medline](#)
- Ikkai A, Curtis CE (2011) Common neural mechanisms supporting spatial working memory, attention and motor intention. *Neuropsychologia* 49:1428–1434. [CrossRef Medline](#)
- Jensen O, Gelfand J, Kounios J, Lisman JE (2002) Oscillations in the alpha band (9–12 Hz) increase with memory load during retention in a short-term memory task. *Cereb Cortex* 12:877–882. [CrossRef Medline](#)
- Kajikawa Y, Schroeder CE (2011) How local is the local field potential? *Neuron* 72:847–858. [CrossRef Medline](#)
- Katzner S, Nauhaus I, Benucci A, Bonin V, Ringach DL, Carandini M (2009) Local origin of field potentials in visual cortex. *Neuron* 61:35–41. [CrossRef Medline](#)
- Lara AH, Wallis JD (2014) Executive control processes underlying multi-item working memory. *Nat Neurosci* 17:876–883. [CrossRef Medline](#)
- Levy RB, Reyes AD (2012) Spatial profile of excitatory and inhibitory synaptic connectivity in mouse primary auditory cortex. *J Neurosci* 32:5609–5619. [CrossRef Medline](#)
- Lundqvist M, Rose J, Herman P, Brincat SL, Buschman TJ, Miller EK (2016) Gamma and beta bursts underlie working memory. *Neuron* 90:152–164. [CrossRef Medline](#)
- Mante V, Sussillo D, Shenoy KV, Newsome WT (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503:78–84. [CrossRef Medline](#)
- Markowitz DA, Curtis CE, Pesaran B (2015) Multiple component networks support working memory in prefrontal cortex. *Proc Natl Acad Sci U S A* 112:11084–11089. [CrossRef Medline](#)
- Mendoza-Halliday D, Torres S, Martinez-Trujillo JC (2014) Sharp emergence of feature-selective sustained activity along the dorsal visual pathway. *Nat Neurosci* 17:1255–1262. [CrossRef Medline](#)
- Meyers EM, Freedman DJ, Kreiman G, Miller EK, Poggio T (2008) Dynamic population coding of category information in inferior temporal and prefrontal cortex. *J Neurophysiol* 100:1407–1419. [CrossRef Medline](#)
- Mongillo G, Barak O, Tsodyks M (2008) Synaptic theory of working memory. *Science* 319:1543–1546. [CrossRef Medline](#)
- Pascanu R, Mikolov T, Bengio Y (2012) On the difficulty of training recurrent neural networks. *Proc 30th Int Conf Mach Learn* 1310–1318.
- Pasternak T, Lui LL, Spinelli PM (2015) Unilateral prefrontal lesions impair memory-guided comparisons of contralateral visual motion. *J Neurosci* 35:7095–7105. [CrossRef Medline](#)
- Patel GH, Shulman GL, Baker JT, Akbudak E, Snyder AZ, Snyder LH, Corbetta M (2010) Topographic organization of macaque area LIP. *Proc Natl Acad Sci U S A* 107:4728–4733. [CrossRef Medline](#)
- Perin R, Berger TK, Markram H (2011) A synaptic organizing principle for cortical neuronal groups. *Proc Natl Acad Sci U S A* 108:5419–5424. [CrossRef Medline](#)

- Rainer G, Asaad WF, Miller EK (1998) Memory fields of neurons in the primate prefrontal cortex. *Proc Natl Acad Sci U S A* 95:15008–15013. [CrossRef Medline](#)
- Rainer G, Rao SC, Miller EK (1999) Prospective coding for objects in primate prefrontal cortex. *J Neurosci* 19:5493–5505. [Medline](#)
- Romo R, Brody CD, Hernández A, Lemus L (1999) Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* 399:470–473. [CrossRef Medline](#)
- Salazar RF, Dotson NM, Bressler SL, Gray CM (2012) Content-specific fronto-parietal synchronization during visual working memory. *Science* 338:1097–1100. [CrossRef Medline](#)
- Sarma A, Masse NY, Wang XJ, Freedman DJ (2016) Task-specific versus generalized mnemonic representations in parietal and prefrontal cortices. *Nat Neurosci* 19:143–149. [CrossRef Medline](#)
- Schall JD, Morel A, Kaas JH (1993) Topography of supplementary eye field afferents to frontal eye field in macaque: implications for mapping between saccade coordinate systems. *Vis Neurosci* 10:385–393. [CrossRef Medline](#)
- Schall JD, Morel A, King DJ, Bullier J (1995) Topography of visual cortex connections with frontal eye field in macaque: convergence and segregation of processing streams. *J Neurosci* 15:4464–4487. [Medline](#)
- Sereno MI, Pitzalis S, Martinez A (2001) Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans. *Science* 294:1350–1354. [CrossRef Medline](#)
- Silver MA, Kastner S (2009) Topographic maps in human frontal and parietal cortex. *Trends Cogn Sci* 13:488–495. [CrossRef Medline](#)
- Song HF, Yang GR, Wang XJ (2016) Training excitatory-inhibitory recurrent neural networks for cognitive tasks: a simple and flexible framework. *PLoS Comput Biol* 12:e1004792. [CrossRef Medline](#)
- Sreenivasan KK, Curtis CE, D’Esposito M (2014) Revisiting the role of persistent neural activity during working memory. *Trends Cogn Sci* 18:82–89. [CrossRef Medline](#)
- Stokes MG (2015) “Activity-silent” working memory in prefrontal cortex: a dynamic coding framework. *Trends Cogn Sci* 19:394–405. [CrossRef Medline](#)
- Stokes MG, Kusunoki M, Sigala N, Nili H, Gaffan D, Duncan J (2013) Dynamic coding for cognitive control in prefrontal cortex. *Neuron* 78:364–375. [CrossRef Medline](#)
- Swaminathan SK, Freedman DJ (2012) Preferential encoding of visual categories in parietal cortex compared with prefrontal cortex. *Nat Neurosci* 15:315–320. [CrossRef Medline](#)
- Swaminathan SK, Masse NY, Freedman DJ (2013) A comparison of lateral and medial intraparietal areas during a visual categorization task. *J Neurosci* 33:13157–13170. [CrossRef Medline](#)
- Swisher JD, Halko MA, Merabet LB, McMains SA, Somers DC (2007) Visual topography of human intraparietal sulcus. *J Neurosci* 27:5326–5337. [CrossRef Medline](#)
- Szatmáry B, Izhikevich EM (2010) Spike-timing theory of working memory. *PLoS Comput Biol* 6:e1000879. [CrossRef Medline](#)
- Wang M, Yang Y, Wang CJ, Gamo NJ, Jin LE, Mazer JA, Morrison JH, Wang XJ, Arnsten AF (2013) NMDA receptors subserve persistent neuronal firing during working memory in dorsolateral prefrontal cortex. *Neuron* 77:736–749. [CrossRef Medline](#)
- Wang XJ (1999) Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *J Neurosci* 19:9587–9603. [Medline](#)
- Zaksas D, Pasternak T (2006) Directional signals in the prefrontal cortex and in area MT during a working memory for visual motion task. *J Neurosci* 26:11726–11742. [CrossRef Medline](#)