HHS Public Access

Author manuscript

Nat Neurosci. Author manuscript; available in PMC 2016 July 01.

Published in final edited form as:

Nat Neurosci. 2016 January ; 19(1): 143–149. doi:10.1038/nn.4168.

Task Specific versus Generalized Mnemonic Representations in Parietal and Prefrontal Cortices

Arup Sarma[1], **Nicolas Y. Masse**[1], **Xiao-Jing Wang**[2,3], and **David J. Freedman**[1,*]

Department of Neurobiology, The University of Chicago, Chicago, IL, 60637, USA

[2]Center for Neural Science, New York University, New York, NY, 10003

[3]NYU-ECNU Institute of Brain and Cognitive Science, NYU Shanghai, Shanghai, China

Abstract

Our ability to learn a wide range of behavioral tasks is essential for responding appropriately to sensory stimuli according to behavioral demands, but the underlying neural mechanism has been rarely examined by neurophysiological recordings in the same subjects across learning. To understand how learning new behavioral tasks impacts underlying neuronal representations, we recorded from posterior parietal cortex (PPC) before and after training on a visual motion categorization task. Here we show that categorization training influenced cognitive encoding in PPC, with a marked enhancement of memory-related delay-period encoding during the categorization task which was absent during a motion discrimination task prior to categorization training. In contrast, the prefrontal cortex (PFC) exhibited strong delay-period encoding during both discrimination and categorization tasks. This reveals a dissociation between PFC's and PPC's roles in working memory, with general engagement of PFC across multiple tasks, in contrast with more task-specific mnemonic encoding in PPC.

Introduction

Humans and other advanced animals have a remarkable ability to perform a wide array of complex behavioral tasks, and to acquire new tasks and skills as a result of learning. Our wide cognitive and behavioral repertoire is essential for making effective decisions, and for adapting to changing behavioral contexts. For example, when interacting with a large group of people, we might need to discriminate between specific individuals in one moment, while the next moment might require categorizing individuals according to their family membership. Neuronal recordings during visual discrimination and categorization tasks have identified encoding of task-related variables across a hierarchy of cortical areas, including task-relevant features in visual cortex, and cognitive factors such as working memory, categorization, and response selection in downstream cortical areas. However, the impact of learning to perform new behavioral tasks on underlying neuronal representations remains

*corresponding author: David J. Freedman, Ph.D., ; Email: dfreedman@uchicago.edu

Author Contributions: D.J.F., A.S., N.Y.M., and X.-J.W. designed the experiments. A.S. and N.Y.M. trained the animals, performed all neurophysiological recordings, performed all data analysis, and assisted in writing the manuscript. D.J.F. supervised the experimental work, assisted in animal training, neurophysiological recordings, data analysis, and wrote the manuscript. X.-J.W. assisted in data analysis and edited the manuscript.

unclear, as few studies have directly compared neuronal encoding in the same animals before and after learning. This kind of information is crucial for understanding the mechanism by which new task-related representations are learned[1].

We examined the role of the posterior parietal cortex (PPC) in mediating visual discrimination and categorization tasks by examining PPC activity before and after categorization task training, and compared visual and mnemonic encoding in PPC with the prefrontal cortex (PFC). The PPC and PFC are both implicated in transforming visual feature encoding in sensory areas into learning-dependent abstract category representations. Previous work showed that both PPC and PFC can encode learned visual motion, spatial and shape categories[2–9], and maintain task-relevant category encoding during delay periods requiring short term memory. In contrast, activity in upstream visual areas such as middle temporal (MT) and inferior temporal cortex (ITC) primarily encodes visual stimulus features, rather than abstract information about their category membership[3,4,10–13].

We recorded from a PPC subregion, the lateral intraparietal (LIP) area, in the same animals before and after learning to perform a visual categorization task. Before categorization training, monkeys were extensively trained on a delayed match to sample (DMS) task in which they had to decide whether a test stimulus was the *same direction* as a previously presented sample. A second recording stage followed long-term training on a delayed match to category (DMC) task in which they indicated whether sample and test stimuli were in the *same category*, defined by a learned category boundary. As expected, categorization training resulted in an emergence of visual category encoding in LIP, consistent with previous studies[14]. However, learning the categorization task also resulted in a striking change in the temporal dynamics of encoding in LIP. During the categorization task, LIP activity showed strong encoding of task-relevant category information during the short-term memory period of the DMC task. In contrast, we did not observe significant delay-period direction encoding in LIP during the DMS task prior to categorization training, even though direction information was task-relevant in that.

Visual working memory is a core executive function, and thought to depend on interactions among a network of cortical and subcortical areas including between parietal and prefrontal cortices[15–18]. A key correlate of working memory in those areas is persistent neuronal activity which selectively encodes task-relevant information[19–21]. The presence of selective delay activity in the DMC but not DMS task suggests that the PPC's role in working memory is highly task-dependent, and raises a question about how remembered motion information was encoded during the DMS task. To address this, we conducted a second study in which we simultaneously recorded from neuronal populations in PPC (including LIP) and PFC during the DMS task. This revealed a sharp contrast in delay activity across parietal and frontal areas, with strong delay-period direction encoding evident only in PFC. This dissociation suggests that PPC plays a limited and task-dependent role in visual working memory, while PFC plays a more generalized role in mnemonic encoding of task relevant information.

Results

Behavioral Tasks: Direction Matching and Categorization

To directly compare neuronal encoding in LIP before and after categorization training, we employed nearly identical behavioral tasks in each recording stage in an effort to equate their behavioral and cognitive demands. The DMS (before categorization training) and DMC (after categorization training) tasks both required monkeys to indicate whether a test stimulus was a “match” to a previous sample by releasing a manual touch-bar. Both tasks used the same timings of task events, and the same eight motion directions were shown as sample stimuli (Fig. 1a). The central difference was that during the DMS task, sample and test stimuli “matched” if they were exactly the same direction (Fig. 1b). In the DMC task (Fig. 1c), sample and test stimuli “matched” if they belonged to the same category, which often meant that matching sample and test stimuli had dissimilar motion directions.

In both tasks, a sample stimulus was followed by a one-second delay and then a test stimulus. If the test matched the sample, the monkey had to release a touch-bar to receive a reward. If the test was a non-match, it was followed by a brief delay and then a matching test stimulus (which required a manual response). The same eight directions were shown as sample stimuli in both tasks, though many more directions were shown during training (see *Methods* and Supplementary Fig. 1). Both monkeys were extensively trained on the DMS task (328 and 243 daily training sessions in monkeys D and H, respectively; see *Methods* and Supplementary Table 1) prior to the start of DMS task neuronal recordings until their discrimination performance reached a stable asymptotic level (Supplementary Fig. 2). During DMS recordings, behavioral performance was greater than chance (50%) when sample and test stimuli were 45° apart, and greater than 85% correct on all other non-match and match conditions (Fig. 2a and Supplementary Fig. 1). During DMC recordings, categorization performance was >85% correct for the four sample directions that were 22.5° from the category boundary, and >90% correct for the four directions that were 67.5° from the boundary (Fig. 2b and Supplementary Fig. 1).

LIP Recordings Before and After Categorization Training

Neuronal recordings were conducted from the same two monkeys during the DMS and DMC recording stages, and targeted overlapping LIP locations in the same hemispheres in each stage. We recorded from 184 LIP neurons (monkey D, N = 92; monkey H, N = 92) during the DMS task, and 270 LIP neurons (monkey D, N = 146; monkey H, N = 124) after training on the DMC task.

During the DMS task, a large fraction of LIP neurons were direction selective (one-way ANOVA on 8 sample directions, $P < 0.01$) during the sample period of the task (N = 115/184 or 62.5%), such as the single neuron examples in Figures 3a–c. During the delay period of the task (excluding the initial 380 ms of the delay; see Methods), the monkeys had to maintain information about the sample direction in short term memory. However, fewer neurons were direction selective during the delay period (N = 21/184 or 11.4%), and those neurons that were direction selective during the delay tended not to show large firing rate

differences between preferred and non-preferred directions, such as the example in Figure 3b (ANOVA, P=0.005).

Following categorization training, we again observed a substantial fraction of direction selective LIP neurons during the sample period of the DMC task (one-way ANOVA, P<0.01, N=182/270 or 67.4%). Interestingly, we found a significantly greater fraction of direction selective neurons during the delay period of the DMC task (N=90/270 or 33.3%) than during the DMS recording stage (chi square test, P<0.01). Notably, among neurons showing delay-period selectivity during the DMC task, most showed selectivity that mirrored the learned categories (e.g. the examples in Figure 3d–f). We examined whether sample-period direction tuning was predictive of delay-period category preferences (i.e. whether the preferred sample direction fell within the preferred category during the delay). Among the population of neurons in the DMC task that were direction selective in the sample period (according to a one-way ANOVA, P<0.01) and category-selective in the delay period (one-way ANOVA, P<0.01), 40/67 LIP neurons (59.7%) of LIP neurons showed corresponding sample-direction and delay-category selectivity.

Temporal Dynamics of Direction and Category Selectivity in LIP

We measured direction and category selectivity at the neuronal population levels using support vector machine (SVM) classifiers applied to pseudopopulations of LIP neurons recorded during each task (see *Methods*). As in a previous study, we employed two specialized SVM classifiers[22] designed to independently assess direction and category selectivity (see *Methods* and Fig. 4a,b). The direction classifier quantified direction selectivity that was independent of category selectivity by examining population activity to the four directions within each category. The category classifier measured category selectivity independent from direction selectivity by training and testing the classifier on pairs of directions in each category with equivalent angular distances. The classifiers were applied in four fixed temporal windows (333 ms width) during the early/late sample and middle/late delay periods. The early delay was not included in order to exclude transient neuronal activity related to sample offset.

The direction classifier revealed direction selectivity in LIP during both the DMS and DMC tasks that was well above chance levels during both early and late sample epochs (Fig. 4c; $P<10^{-4}$, bootstrap, see *Methods*), and did not vary significantly between the two tasks (P>0.05, bootstrap, see *Methods*). However during the delay period, the direction classifier was unable to robustly decode direction information during the DMS task, even though the monkeys had to remember the sample direction during the delay period. For example, we did not observe direction classification that was significantly above chance levels in the middle delay epoch (Fig. 4c; P=0.16, bootstrap), although it did reach significance in the late delay (P=0.006, bootstrap). During the DMC task, middle and late delay period direction decoding performance was modestly elevated above the chance (0.25) level, and reached significance in both delay epochs (Fig. 4c; middle: P=0.004, late: P=0.01). However, it was still considerably weaker than during the sample epochs.

In contrast to the direction classifier results, the category classifier showed significant category decoding during both middle and late delay epochs of the DMC task (Fig. 4d;

middle: $P<10^{-4}$, late: $P<10^{-4}$, bootstrap), which was also evident during the late (but not early) sample (Fig. 4d; P=0.001). We did not expect to observe category selectivity during the DMS task, prior to categorization training. This was confirmed, as significant category selectivity was not observed during the DMS task during the sample (Fig. 4d; early: P=0.72, late: P=0.53) or delay (Fig. 4d; middle: P=0.60, late: P=0.94). These results suggest that categorization training had a marked impact on the time-course of neuronal encoding in LIP, with a training related enhancement of selectivity during the delay period.

To characterize the time-course of direction and category encoding in LIP in more detail, we applied both classifiers using a sliding window (width=200 ms, step size=10 ms) in which the classifiers were trained and tested at each time step (Fig. 5a,b). Statistical significance was assessed at each time point, and significant time bins are indicated by the light (P<0.05) or dark (P<0.01) colored tick marks at the top or bottom of each plot (see *Methods*). For direction selectivity, we found a striking similarity in the strength and time-course of selectivity in the DMS and DMC tasks, with strong direction decoding during the sample period which returned to near-chance levels during the delay period (Fig. 5a). The category classifier showed an emergence of category selectivity during the late sample period of the DMC task which was maintained through the delay (Fig. 5b). As expected, category selectivity during the DMS task remained at near-chance levels throughout the trial (Fig. 5b).

The observed patterns of direction and category selectivity during the DMS and DMC tasks did not depend on this specific analysis approach (SVM classification) used. For example, we found qualitatively similar results regarding the relative strength and time-course of selectivity when examining direction and category selectivity using selectivity indices applied to individual LIP neurons (Supplementary Fig. 3, and *Methods*).

A number of individual LIP neurons exhibited stable category selectivity throughout the delay period of the DMC task (Fig. 3d–f), similar to that observed at the population level (Fig. 5b). We assessed the stability of neuronal category selectivity during the DMC task at the LIP population level by training and testing the category classifier at different time points in the trial. For example, if a classifier trained using neuronal data from the beginning of the trial shows high decoding performance when tested at later time points, this indicates stable patterns of neuronal encoding across the population between those time points. This revealed a large and stable period of high category classification accuracy spanning the entire delay period, as shown by the uniform regions of high accuracy in Figure 5c (right panel). As expected, when applied to LIP data from the DMS task prior to categorization training, the category classifier produced near-chance decoding accuracies for all combinations of training and testing times during the delay (Fig. 5c, left panel). We used the same approach to examine the temporal stability of LIP direction selectivity in the DMS and DMC tasks. This revealed strong and relatively stable direction selectivity during sample presentation, but little direction selectivity (neither stable nor dynamic) during the delay (Supplementary Fig. 4).

An important issue is to understand whether delay-period encoding in LIP in the DMC but not DMS tasks was due to the difference between the tasks themselves, or rather to

differences in the duration of behavioral training leading up to recordings in each tasks. We addressed this concern in several ways, and found that the appearance of delay-period encoding in the DMC task was not a result of training duration, suggesting that delay-period encoding was instead task dependent. First, the monkeys were trained extensively on the DMS task prior to DMS-task recordings (for ~1 year of daily training sessions; Supplementary Table 1) until they reached a stable level of task performance (Supplementary Fig. 2). The monkeys' extensive training and apparent plateau in their level of task-performance makes it is unlikely that further DMS training would have markedly affected their behavioral performance or underlying neuronal encoding.

We also addressed this concern by examining additional LIP data from 184 neurons (monkey D, N = 70; monkey H, N = 114) collected at an intermediate training stage between the DMS and fully-trained DMC tasks (104 and 70 DMC training sessions in monkeys D and H, respectively, following DMS recordings; see *Methods*). At the mid-training DMC stage, the monkeys performed the DMC task at greater than chance levels, though below the level of performance in the fully-trained DMC task (Supplementary Fig. 5a). Furthermore, their behavioral performance indicates that they were at least partially using a similarity strategy (based on the difference between sample and test directions) and had not learned the abstract groupings between directions in each category (Supplementary Figs. 5b–e). During the mid-training DMC stage, category and direction encoding were both evident during the sample period. Importantly, LIP activity in this stage showed a lack of significant delay-period category or direction encoding (Supplementary Fig. 6). This confirms that the emergence of LIP delay-period encoding in the final DMC stage is not due to training duration, as we would instead have expected a gradual emergence of delay encoding throughout training rather than an abrupt emergence at the final training stage.

Comparison of Visual and Mnemonic Encoding in PPC and PFC

The observed lack of strong delay period direction selectivity in LIP during the DMS task raises questions about how direction information was stored in short-term memory, and whether other cortical areas such as the PFC play a more general role in visual working memory. To address this, we simultaneously recorded from ensembles of neurons in the PPC (including LIP; monkey Q, N = 62; monkey W, N = 93) and lateral PFC (monkey Q, N = 127; monkey W, N = 131) of two monkeys (different animals than in previous studies) during the direction DMS task. The DMS task was nearly identical to that in Figure 1a,b, except that a different set of sample and test directions were used (see *Methods*). The monkeys performed the DMS task with high average accuracy during neuronal recordings (monkey Q: 91.5% correct; monkey W: 89.1% correct; Supplementary Fig. 7).

Similar to our previous LIP results, many PPC neurons were direction selective (one-way ANOVA on 8 sample directions, $P < 0.01$) for at least one out of two receptive field (RF) placements (see Methods) during the sample period (76/155 neurons), but many fewer neurons were selective during the delay period (26/155), such as the example PPC neuron shown in Figure 6a. In contrast, many PFC neurons, such as the example shown in Figure 6b, were direction selective during both the sample (83/258) and delay (75/258) periods.

To compare direction selectivity in PPC and PFC, we applied the same direction classifier (as above) to the PPC and PFC data during the DMS task. The observed results in PPC replicated our findings in LIP described above—a lack of significant delay-period direction selectivity, but significant direction selectivity during the sample period. In contrast to the lack of strong delay period selectivity in PPC, the direction classifier revealed significant direction selectivity in PFC during the delay as well as the sample (bootstrap, $P<0.01$). Thus PFC showed a preferential delay-period encoding of task-relevant direction information compared to PPC during the DMS task, suggesting that it plays a more generalized role in visual working memory than PPC. We note that direction decoding accuracy in PPC (Fig. 7a) is weaker than that observed in the LIP study described above (Fig. 5a). This difference is likely due to the broader sampling of PPC in these experiments (which likely included both LIP and neighboring parietal areas such as the medial intraparietal area and 7a), or the inability to place the stimulus in the RF of all recorded neurons when using multielectrode arrays (compared to a single electrode in the earlier studies; see Methods).

We repeated the classifier approach used above to assess stable vs. dynamic direction encoding in PPC and PFC by training and testing the classifier at different time points. In PPC this revealed strong direction selectivity in the sample period and neither stable nor dynamic selectivity in the delay (Supplementary Fig. 8). PFC showed an apparent mixture of stable and dynamic direction selectivity in both the sample and delay, indicated by the appearance of rectangular regions of elevated classification accuracy (i.e. stable encoding) during both time periods, with a superimposed diagonal zone of highest classification values (i.e. dynamic encoding) corresponding to when the classifier was trained and tested at the same time points (Fig. 7b). This is consistent with previous studies which found stable[16,19] and/or dynamic[23–25] delay period selectivity in PFC. A final note is that, although we do not have PPC and PFC data from these two animals during the fully-trained DMC task, a previous study from our group examining PFC during the fully trained DMC task found strong and temporally stable category selectivity[9].

Discussion

Taken together, these results give insight into the task dependence and plasticity of visual and mnemonic encoding in PPC, and reveal distinct roles of PPC and PFC in visual working memory. First, long-term categorization training produced striking changes in the type of information encoded in PPC, and in the time-course of that encoding. Before categorization training, PPC showed robust direction tuning only during stimulus presentation, but not during the delay period when direction information had to be maintained in short term memory. In contrast, PPC recordings after categorization training revealed strong and stable category encoding throughout the memory delay period. Second, the comparison of PPC and PFC during the DMS task reveals that PFC plays a more generalized role in working memory than PPC. During the DMS and DMC tasks, delay period activity in PFC encodes task-relevant direction or category information during the DMS or DMC tasks, respectively. This contrasted with task-specific delay activity in PPC, which showed working memory encoding only during the categorization task.

The observation that visual categorization training affected visual feature encoding in LIP is consistent with previous work showing that PPC activity can reflect learned categories[3,5,6,8,9,14,20] and stimulus-stimulus associations[8], while MT encoding appears not to be significantly influenced by long-term categorization training[3,26]. Categorization training also had a marked and unexpected impact on the time-course of selectivity in LIP throughout the trial, with the appearance of stable delay period selectivity during the DMC but not DMS task. The two tasks used identical sample stimuli, timings of task events, motor responses, and requirement to remember sample information during the delay. Thus, the presence of delay encoding in the categorization task arose as a result of learning the DMC task and its specific task demands, such as the requirement to convert encoding of stimulus direction into an abstract categorical variable. This is also consistent with LIP playing a more direct role in mediating performance of the DMC than DMS task[9,14,27,28].

The comparison of PFC and PPC during the DMS task revealed that PFC, but not PPC, robustly encoded task-relevant direction information during the memory delay period. Furthermore, previous work from our group and others found that PFC also shows delay period encoding of task-relevant information during visual categorization tasks[2–9], in addition to a variety of other working memory tasks[18,19]. This supports a generalized role for PFC in working memory both during tasks which require mnemonic encoding of stimulus features (as in the motion DMS task)[29] or more abstract cognitive variables, such as a categorical decision. This contrasts with the lack of significant delay-period encoding in PPC during the DMS task, suggesting that PPC is selectively engaged by non-spatial working memory tasks with greater cognitive demands or abstraction[30,31]. Interestingly, strong delay period encoding is evident in both PPC and PFC during relatively simple spatial delayed response tasks[32], suggesting that non-spatial and spatial working memory rely on distinct underlying mechanisms. Furthermore, we previously found that both PFC and LIP showed strong delay-period category encoding during the DMC task[9], suggesting that the two areas are more coordinated during categorization than discrimination tasks.

In contrast to the task-dependence of LIP delay encoding, sample period activity was strongly direction selective both before and after categorization training. This is consistent with previous reports of LIP direction tuning during both passive viewing and complex tasks[33,34], and likely reflects LIP's interconnection with upstream motion processing areas such as MT, the medial superior temporal (MST)[35], and the ventral intraparietal[36–38]. Interestingly a recent study using a variant of a motion direction DMS task found delay-period direction encoding in MST but not MT[39]. The observation of direction selective delay activity in both MST and PFC raises the possibility that interactions between these two areas (but not MT or LIP) may underlie working memory for direction. An additional open question is whether the strong sample-period direction encoding in LIP during both DMS and DMC tasks is related to the monkeys' extensive motion discrimination training prior recordings in the DMS task. This difference in training history might also account for the apparently stronger direction relative to category encoding[20] as well as the more gradual emergence of neuronal category encoding[8] during the sample period in this compared to previous reports from our group,.

It will be important to extend this work to examine the real-time impact of categorization training during the learning process itself. Neuronal correlates of rapid within-session associative or category learning have been demonstrated in PFC, striatum, and inferior temporal cortex[40–43] and recordings during long term training on a perceptual decision task revealed learning-related enhancement of motion encoding in LIP but not MT[26]. A related issue is to understand neuronal mechanisms underlying the ability to rapidly switch between tasks. Previous work found that PFC activity can reflect rapidly changing category rules[44,45], although it remains to be seen whether DMS-DMC task switching is accompanied by rapid changes in delay-period encoding. Alternatively changes in the dynamics of delay-period encoding may have developed gradually during categorization training.

Online Methods

The detailed methods used in this study have been described in previously[22,28], and are briefly described below.

Behavioral Tasks and Stimulus Display

Monkeys performed delayed match-to-sample (DMS) and delayed-match-to-category (DMC) tasks using 360° of motion directions as sample and test stimuli. Monkeys released a manual lever to indicate whether a test stimulus was the same direction (DMS) or same category (DMC) as a previously presented sample. Monkeys were required to maintain fixation within 2.0°–2.5° of a 0.2° fixation point throughout the entire trial, and stimuli were always presented in LIP neurons' RFs. For the monkeys D and H in the first experiment (Figs. 1–5), sample and test stimuli were 9.0° diameter circular patches of high-contrast, 100% coherent random dots. Dots moved at $12°s^{-1}$ and were displayed at a frame rate of 75 Hz. For monkeys Q and W in the second experiment (Figs. 6 and 7), stimuli had a diameter of 6.0° and moved at $12°s^{-1}$. Identical task timings and rewards were using during all tasks. During DMS-task recordings, we used 8 evenly-spaced sample directions were used as sample stimuli. During DMS task recordings, sample and non-matching test stimuli were separated by either 45°, 60°, or 75° in order to keep monkeys' performance on the task above chance but below complete certainty. The sets of test stimulus directions in the DMS and DMC tasks differed slightly, as the DMS set included additional motion directions (Fig. 1a and Methods). During DMC-task recordings, test stimuli were chosen randomly from the same directions as the sample stimuli. However, this study focuses on encoding related to sample stimuli and does not focus on the test-period.

Stimulus presentation, task events, rewards, and behavioral data acquisition were accomplished using an Intel-based PC equipped with MonkeyLogic software running in MATLAB[46]. Gaze positions were measured and recorded at a sampling rate of 1.0 kHz using an EyeLink 1000 optical eye tracker (SR Research). Visual stimuli were presented on a 21" color CRT monitor (1,280 × 1,024 resolution, 75 Hz refresh rate, 47-cm viewing distance).

Behavioral Training

Monkeys D and H (experiment 1) were each trained on the motion DMS task for approximately one year (~1400 correct trials/session; Supplementary Table 1 and Supplementary Fig. 2) prior to DMS recordings. Monkeys were trained each day on 12 unique sample directions, and test directions were chosen to include a wide range of angular differences between non-matching sample and test directions. Early DMS training sessions focused on large angular differences between sample and test directions (e.g. 90°–180°) until monkeys' accuracy reached a high level. Middle to late stage training sessions focused on more challenging sample-test differences (e.g.5°–30°), and DMS training continued until the monkeys' performance improved and stabilized. Monkeys Q and W (experiment 2) were trained on the DMS task with 6 evenly-spaced sample directions for >150 training sessions. These same directions were used for recording in experiment 2. Test stimuli were either an identical direction match, or were 180° from the sample direction.

DMC task training of Monkeys D and H was conducted in two stages. The first was the mid-training DMC stage (104 and 70 sessions in monkeys D and H, respectively; Supplementary Table 1) in which the category boundary was introduced and the monkeys were rewarded for indicating (with a lever release) whether the test stimulus was in the same category as the sample. In this training stage, the 12 sample directions used for training were shown with equal frequency during each session, and a wide range of sample-test differences were used. At the end of this training stage (defined by the monkeys reaching stable level of performance), the mid-training DMC recordings were conducted from LIP. In the second DMC training stage (78 and 65 sessions in monkeys D and H; Supplementary Table 1), we over-emphasized near-boundary (i.e. 15° from boundary) sample stimuli for which the monkeys had shown lower performance in the middle training stage. This training stage continued until the monkeys' performance for near-boundary stimuli improved and stabilized (Fig. 2b and Supplementary Fig. 5).

Electrophysiological Recording

Recording chambers and headposts were implanted on all four male monkeys (Macaca mulatta, 9.0–12.0 kg, 7–9 years old). For monkeys D and H, chambers were placed over the intraparietal sulcus (IPS) as determined by stereotaxic coordinates obtained from MRI scans prior to implantation of the headpost and recording chamber. Recording chambers were centered ~1.0 mm posterior to the intra-aural line and ~13.0 mm lateral from the midline. Recordings were conducted using single 75 μm tungsten microelectrodes (FHC), a dura piercing guide tube, and a motorized microdrive system (NAN instruments).

For monkeys Q and W, two 32-channel semi-chronic recording systems (Gray Matter Research, Bozeman, MT) were used to record from PPC and PFC. MRI scans were used to guide chamber placement. For PPC recordings, chambers were placed over the IPS, ~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. We advanced all PPC electrodes until their estimated positions were below the IPS, guided by its known anatomical depth. Additional evidence for electrode depth on many recording channels was the marked reduction in spiking activity as electrodes entered

the sulcus. 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 μm tungsten microelectrodes (Alpha-Omega). Before each session, we lowered electrodes between 0 and 1 mm in order to optimally record the spiking activity of well isolated neurons.

All surgical and experimental procedures followed the University of Chicago's Animal Care and Use Committee and US National Institutes of Health guidelines. Monkeys were housed in individual cages under a 12-hour 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 (Plexon Inc.) to verify the quality and stability of neuronal isolations.

Receptive Field Mapping and Stimulus Placement

For monkeys D and H, neurons were tested with a memory-saccade task prior to DMS or DMC recordings in order to identify area LIP and determine neurons' RF locations. Neurons were considered to be in LIP if they showed spatially selective visual, delay, and/or presaccadic activity during the memory-saccade task or were located between such neurons. LIP neurons were not pre-screened for direction or category selectivity or responsiveness to visual motion stimuli. The range of eccentricity for saccade amplitudes tested and stimulus placement ranged from 7.0°–10.0°. Neuronal responses during the memory-saccade task were analyzed in real time by automated MATLAB scripts that displayed neuronal spatial selectivity. Sample and test stimuli during the DMS and DMC tasks were placed within LIP RFs.

For monkeys Q and W, we recorded from all neurons with well isolated action potentials. In the PFC-PPC simultaneous recordings, we could not place stimuli within the RFs of all recorded neurons. To increase the chances that stimuli were in or near neuronal RFs, we ran the experiment in alternating 10-trial blocks in which stimulus position was varied between two non-overlapping positions (7.0° eccentricity; ±45° relative to horizontal meridian) in the visual field that was contralateral to the hemisphere targeted for neuronal recordings. Analysis of neuronal data revealed qualitatively similar results (in both cortical areas) for each of the two stimulus locations considered separately, or when trials from the two locations were combined. Thus we combined trials for both stimulus locations in the manuscript.

Data Analysis

All analyses were conducted on correct trials. The patterns of behavioral and neuronal results were similar, and all main effects were observed in both monkeys. Thus, the two datasets from each animal in each study were combined for all population analyses.

Whole-epoch based analyses were computed for each neuron in two epochs. The sample period was a 667 ms epoch beginning 80 ms after sample stimulus onset (to account for neuronal response latency). The delay period was a 713 ms epoch beginning 380 ms after stimulus offset (to exclude firing rates shortly after stimulus offset). Split epoch analyses

were computed on four epochs, each 333 ms long. Early and late sample periods were consecutive 333 ms epochs beginning 80 ms after stimulus onset. Middle and late delay periods were consecutive 333 ms epochs beginning 427 ms after stimulus offset.

A population decoding approach was used to quantify direction and category information at the population level. We trained and tested linear classifiers on surrogate populations of neurons constructed from the individual neuron recordings[22,47]. When we compared classification accuracy between two populations, we always made sure to use an equal number of neurons when classifying each population. Specifically, if the population count of the two populations was n1 and n2, where n1<=n2, we would bootstrap n1 neurons from each population to use for the classifier. We would select a different set of neurons for each of the 10,000 iterations for the partial-epoch, or for each of the 100 iterations for the sliding window analysis. During each of the 10,000 repeats, we selected a different set of neurons (using the procedure described above), and a different set of 20 trials for each sample direction (new page 13). Since classification scores will depend both on which neurons and which trials are included, randomly selecting both fully captures the variance in our model. It should be noted that randomly selecting neurons and trials on each iteration makes it more difficult to obtain significance, compared to selecting trials or neurons alone.

In order to measure category selectivity separately from direction selectivity, different classifiers were constructed for each set of two adjacent directions within a category and their diametrically (and hence categorically) opposite sample directions. These classifiers were trained on the pool of trials from the remaining two sample directions in each category, which (on average) have the same distance from the directions in the testing set, and then tested on the two sets of opposing directions in the test set. Training and testing sets were constructed by randomly sampling 20 trials for each sample direction. Neurons were randomly sampled from the total population of neurons in each task in order to compare accuracy between surrogate neuronal populations of the same size. During partial-epoch analyses, this was repeated 10,000 times for each epoch. During sliding window analyses, this was repeated 100 times for each 10 ms time point. Overall performance was found by averaging performance across the two classifiers at each time point. By training and testing on separate sample directions, this classifier indicates, at a population level, whether firing rates tend to be similar within a category and different between categories. Therefore, a population of Gaussian-tuned neurons with uniformly distributed preferred directions would produce chance performance on this classifier (which we confirmed using simulated data).

Similarly, we quantified direction selectivity separately from category selectivity by constructing separate classifiers for trials with directions within each category. For the directions within each category, we randomly sampled 20 trials for each sample direction. Neurons were randomly sampled from the total population in each task in order to compare accuracy between surrogate neuronal populations of the same size. 75% of these trials were used for training the classifier, and 25% were used for testing using a 4-fold cross-validation approach. During partial-epoch analyses, this was repeated 10,000 times for each epoch. During sliding window analyses, this was repeated 100 times for each 10 ms time point. Overall performance was found by averaging performance between the two classifiers at each time point. A neuron with identical firing rates within each category would not

contribute to the performance of this classifier. Both direction and category classifiers were trained using a support vector machine. The direction classifier was trained using multi-class SVM[48]. Spike counts were normalized between −1 and +1.

The stability of direction and category selectivity was quantified by extending the population decoding approach to compare patterns of information coding throughout the trial. Rather than training and testing the classifiers using data from the same time point in the trial, the stability of encoding was assessed by training at one time point and then testing on data from every time point in the trial. Since the training and testing directions are not independent for the direction classifier, one potential concern with this analysis was overlapping training and testing sets of trials between time points. Accordingly, for each repetition of training the classifier at a given time point, every testing time point used the same, separate set of testing trials (with firing rates at different times). Importantly, these testing trials were always different from the training trials used, and hence there was no overlap in spikes between training and testing sets. This was not a concern for the category classifier, since training and testing directions were independent. In both cases, high classifier accuracy at different training and testing times indicates that stimulus information can be decoded using the same pattern of firing rates at both times, while chance or below chance accuracy indicate otherwise. Thus, contiguous blocks of high classifier accuracy across time suggest that stable representations of stimulus information are being used for the population code.

Category selectivity was quantified for individual neurons using a receiver operating characteristic (ROC)-based category tuning index (rCTI) that compared differences in the distributions of firing rates across trials within and between categories[22,28]. We defined the between-category discrimination (BCD) as the average absolute value of the area under the ROC curve subtracted by 0.5 for the distributions of firing rates for all pairs of sample directions where directions were in different categories. This meant that, for example, the ROC values for pairs of directions which were 45° apart were weighted three times as heavily as pairs of directions that were 135° apart, since there were three times as many pairs of directions that were separated by 135°. The within-category discrimination (WCD) was defined similarly, but using pairs of directions within the same category. Since pairs of directions separated by 180° only existed between categories, these pairs were not included in the rCTI computation. The rCTI was then defined as rCTI = BCD – WCD. Values of the rCTI can range from +0.5 (perfect discrimination between categories, identical firing rates within categories) to −0.5 (large differences in firing rates within categories, no differences between categories). Values of zero indicated equal discriminability for firing rates within and between categories.

Direction selectivity was quantified for individual neurons using an ROC-based direction selectivity index (rDSI) which measured the maximum discriminability of firing rates between pairs of sample directions within the same category. This calculation involved 100 repetitions, in which we randomly selected with replacement 20 trials per motion direction for each repetition. For each repetition, the absolute value of the area under the ROC curve subtracted by 0.5 was computed for the firing rates from the sets of trials for each pair of sample directions within the same category. The maximum value was chosen for each

repetition. The rectification of ROC values around 0.5 meant that all variability contributed positively to ROC values, and hence values were positively biased. We subtracted this bias per repetition by performing the same computation for trials with sample direction labels shuffled and subtracting the shuffled value from the correctly labeled value. The final rDSI value was the average of these shuffle-corrected values over all repetitions. Values of the rDSI could theoretically range from −0.5 (large differences in firing for shuffled label trials, no difference in firing for correctly labeled trials) to +0.5 (large differences in firing for correctly labeled trials, no difference in firing for shuffled labeled trials), although in practice the lower bound was close to zero since shuffled label trials had minimal differences in firing rates.

Statistics

No statistical methods were used to pre-determine sample sizes but our sample sizes are similar to those generally employed in the field. The criteria used to select neurons for recording in LIP (spatially-selective activity on the memory-saccade task) and PPC/PFC (all neurons with well isolated action potentials) necessarily meant that sample selection was not randomized or blinded.

For the purpose of comparing whole-epoch direction selectivity between DMS and DMC tasks, direction selectivity was determined for individual neurons via significance on one-way ANOVA (eight sample directions, $P < 0.01$). The statistical significance of the difference in proportions of direction selective neurons between DMS and DMC tasks was computed using a chi square test. Data distribution was assumed to be normal but this was not formally tested.

The statistical significance of classification accuracies above chance were calculated using a shuffle analysis. A null distribution was produced by training and testing the classifiers described above on the same population of data with trial sample direction labels shuffled. Average classification accuracy was considered significantly greater than chance if the value was greater than 95% ($p<0.05$) or 99% ($p<0.01$) of the values from the null distribution. Since we only tested for classification accuracies significantly greater than the shuffled null distribution, this was equivalent to a one-tailed test.

The statistical significance of differences in classification accuracy between the DMS and DMC tasks was calculated using a bootstrap analysis. The classification accuracy values from each repetition of the classifier at each epoch or time point were compared between DMS and DMC populations. For partial epoch-based analyses, this totaled 10,000×10,000 = 1×10^8 comparisons per epoch; for sliding window analyses, this totaled 100×100 = 10,000 comparisons per time point. Average classification accuracy was deemed to be significantly different between the DMS and DMC tasks if 97.5% ($p<0.05$) or 99.5% ($p<0.01$) of the values from one task were greater than those from the other. Since we were testing for positive or negative classification accuracy differences between tasks, this was equivalent to a two-tailed test.

The statistical significance of differences in category selectivity (rCTI) and direction selectivity (rDSI) between the populations of neurons recorded during the DMS and DMC tasks was computed using a two-sided Wilcoxon rank-sum test.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank Navaneethan Santhanam, Rutwik Kharkar, and Maimon Rose for assistance with animal training, Jonathan Hodnefield for expert behavioral and experimental assistance, Stephanie Thomas for expert technical assistance, and staff of The University of Chicago Animal Resources Center for expert veterinary assistance. We also thank the following for valuable input and/or comments on an earlier version of this manuscript: John Assad, Sliman Bensmaia, Nicolas Brunel, Jah Chaisangmongkon, Tatiana Engel, Guilhem Ibos and Krithika Mohan. This work was supported by a NIH R01MH092927 (X.J.W. and D.J.F.), NIH R01EY019041 (D.J.F), and NRSA MH097428 (A.S.). Additional support was provided by a McKnight Scholar award, the Alfred P. Sloan Foundation, NSF CAREER award 0955640, The Brain Research Foundation (D.J.F.), and R01-MH062349 (X.-J.W.).

References

1. Engel TA, Chaisangmongkon W, Freedman DJ, Wang X-J. Choice-correlated activity fluctuations underlie learning of neuronal category representation. Nat. Commun. 2015; 6:6454. [PubMed: 25759251]
2. Freedman DJ, Riesenhuber M, Poggio T, Miller EK. Categorical representation of visual stimuli in the primate prefrontal cortex. Science. 2001; 291:312–316. [PubMed: 11209083]
3. Freedman DJ, Assad JA. Experience-dependent representation of visual categories in parietal cortex. Nature. 2006; 443:85–88. [PubMed: 16936716]
4. Freedman DJ, Riesenhuber M, Poggio T, Miller EK. A comparison of primate prefrontal and inferior temporal cortices during visual categorization. J. Neurosci. 2003; 23:5235–5246. [PubMed: 12832548]
5. Goodwin SJ, Blackman RK, Sakellaridi S, Chafee MV. Executive control over cognition: stronger and earlier rule-based modulation of spatial category signals in prefrontal cortex relative to parietal cortex. J. Neurosci. 2012; 32:3499–3515. [PubMed: 22399773]
6. Crowe DA, et al. Prefrontal neurons transmit signals to parietal neurons that reflect executive control of cognition. Nat. Neurosci. 2013; 16:1484–1491. [PubMed: 23995071]
7. Ferrera VP, Yanike M, Cassanello C. Frontal eye field neurons signal changes in decision criteria. Nat. Neurosci. 2009; 12:1458–1462. [PubMed: 19855389]
8. Fitzgerald JK, Freedman DJ, Assad JA. Generalized associative representations in parietal cortex. Nat. Neurosci. 2011; 14:1075–1079. [PubMed: 21765425]
9. Swaminathan SK, Freedman DJ. Preferential encoding of visual categories in parietal cortex compared with prefrontal cortex. Nat. Neurosci. 2012; 15:315–320. [PubMed: 22246435]
10. Freedman DJ, Riesenhuber M, Poggio T, Miller EK. Experience-dependent sharpening of visual shape selectivity in inferior temporal cortex. Cereb. Cortex. 2006; 16:1631–1644. [PubMed: 16400159]
11. Sigala N, Logothetis NK. Visual categorization shapes feature selectivity in the primate temporal cortex. Nature. 2002; 415:318–320. [PubMed: 11797008]
12. Vogels R. Categorization of complex visual images by rhesus monkeys. Part 2: single-cell study. Eur. J. Neurosci. 1999; 11:1239–1255. [PubMed: 10103119]
13. De Baene W, Ons B, Wagemans J, Vogels R. Effects of category learning on the stimulus selectivity of macaque inferior temporal neurons. Learn. Mem. 2008; 15:717–727. [PubMed: 18772261]
14. Fitzgerald JK, Swaminathan SK, Freedman DJ. Visual categorization and the parietal cortex. Front. Integr. Neurosci. 2012; 6:18. [PubMed: 22582040]

15. D'Esposito M, Postle BR. The Cognitive Neuroscience of Working Memory. Annu. Rev. Psychol. 2014
16. Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci. 2001; 24:167–202. [PubMed: 11283309]
17. Curtis CE. Prefrontal and parietal contributions to spatial working memory. Neuroscience. 2006; 139:173–180. [PubMed: 16326021]
18. Lara AH, Wallis JD. Executive control processes underlying multi-item working memory. Nat. Neurosci. 2014; 17:876–883. [PubMed: 24747574]
19. Goldman-Rakic PS. Cellular basis of working memory. Neuron. 1995; 14:477–485. [PubMed: 7695894]
20. Sreenivasan KK, Curtis CE, D'Esposito M. Revisiting the role of persistent neural activity during working memory. Trends Cogn. Sci. 2014; 18:82–89. [PubMed: 24439529]
21. Wang X-J. Synaptic reverberation underlying mnemonic persistent activity. Trends Neurosci. 2001; 24:455–463. [PubMed: 11476885]
22. Swaminathan SK, Masse NY, Freedman DJ. A comparison of lateral and medial intraparietal areas during a visual categorization task. J. Neurosci. 2013; 33:13157–13170. [PubMed: 23926269]
23. Crowe DA, Averbeck BB, Chafee MV. Rapid Sequences of Population Activity Patterns Dynamically Encode Task-Critical Spatial Information in Parietal Cortex. J. Neurosci. 2010; 30:11640–11653. [PubMed: 20810885]
24. Meyers EM, Freedman DJ, Kreiman G, Miller EK, Poggio T. Dynamic population coding of category information in inferior temporal and prefrontal cortex. J. Neurophysiol. 2008; 100:1407–1419. [PubMed: 18562555]
25. Zaksas D, Pasternak T. Directional Signals in the Prefrontal Cortex and in Area MT during a Working Memory for Visual Motion Task. J. Neurosci. 2006; 26:11726–11742. [PubMed: 17093094]
26. Law C-T, Gold JI. Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nat. Neurosci. 2008; 11:505–513. [PubMed: 18327253]
27. Freedman DJ, Assad JA. A proposed common neural mechanism for categorization and perceptual decisions. Nat. Neurosci. 2011; 14:143–146. [PubMed: 21270782]
28. Rishel CA, Huang G, Freedman DJ. Independent category and spatial encoding in parietal cortex. Neuron. 2013; 77:969–979. [PubMed: 23473325]
29. Pasternak T, Lui LL, Spinelli PM. Unilateral prefrontal lesions impair memory-guided comparisons of contralateral visual motion. J. Neurosci. 2015; 35:7095–7105. [PubMed: 25948260]
30. Koenigs M, Barbey AK, Postle BR, Grafman J. Superior Parietal Cortex Is Critical for the Manipulation of Information in Working Memory. J. Neurosci. 2009; 29:14980–14986. [PubMed: 19940193]
31. Berryhill ME, Chein J, Olson IR. At the intersection of attention and memory: the mechanistic role of the posterior parietal lobe in working memory. Neuropsychologia. 2011; 49:1306–1315. [PubMed: 21345344]
32. Chafee MV, Goldman-Rakic PS. Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. J. Neurophysiol. 1998; 79:2919–2940. [PubMed: 9636098]
33. Fanini A, Assad JA. Direction selectivity of neurons in the macaque lateral intraparietal area. J. Neurophysiol. 2009; 101:289–305. [PubMed: 18987126]
34. Ibos G, Freedman DJ. Dynamic Integration of Task-Relevant Visual Features in Posterior Parietal Cortex. Neuron. 2014; 83:1468–1480. [PubMed: 25199703]
35. Born RT, Bradley DC. Structure and function of visual area MT. Annu. Rev. Neurosci. 2005; 28:157–189. [PubMed: 16022593]
36. Colby CL, Duhamel JR, Goldberg ME. Ventral intraparietal area of the macaque: anatomic location and visual response properties. J. Neurophysiol. 1993; 69:902–914. [PubMed: 8385201]

37. Cook EP, Maunsell JHR. Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey. J. Neurosci. 2002; 22:1994–2004. [PubMed: 11880530]

38. Lewis JW, Van Essen DC. Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. J. Comp. Neurol. 2000; 428:112–137. [PubMed: 11058227]

39. Mendoza-Halliday D, Torres S, Martinez-Trujillo JC. Sharp emergence of feature-selective sustained activity along the dorsal visual pathway. Nat. Neurosci. 2014; 17:1255–1262. [PubMed: 25108910]

40. Antzoulatos EG, Miller EK. Differences between neural activity in prefrontal cortex and striatum during learning of novel abstract categories. Neuron. 2011; 71:243–249. [PubMed: 21791284]

41. Asaad WF, Rainer G, Miller EK. Neural activity in the primate prefrontal cortex during associative learning. Neuron. 1998; 21:1399–1407. [PubMed: 9883732]

42. Pasupathy A, Miller EK. Different time courses of learning-related activity in the prefrontal cortex and striatum. Nature. 2005; 433:873–876. [PubMed: 15729344]

43. Messinger A, Squire LR, Zola SM, Albright TD. Neuronal representations of stimulus associations develop in the temporal lobe during learning. Proc. Natl. Acad. Sci. U. S. A. 2001; 98:12239–12244. [PubMed: 11572946]

44. Cromer JA, Roy JE, Miller EK. Representation of Multiple, Independent Categories in the Primate Prefrontal Cortex. Neuron. 2010; 66:796–807. [PubMed: 20547135]

45. Roy JE, Riesenhuber M, Poggio T, Miller EK. Prefrontal Cortex Activity during Flexible Categorization. J. Neurosci. 2010; 30:8519–8528. [PubMed: 20573899]

Methods-only References

46. Asaad WF, Santhanam N, McClellan S, Freedman DJ. High-performance execution of psychophysical tasks with complex visual stimuli in MATLAB. J. Neurophysiol. 2013; 109:249–260. [PubMed: 23034363]

47. Meyers EM, Qi X-L, Constantinidis C. Incorporation of new information into prefrontal cortical activity after learning working memory tasks. Proc. Natl. Acad. Sci. U. S. A. 2012; 109:4651–4656. [PubMed: 22392988]

48. Chang C, Lin C. LIBSVM: a library for support vector machines. ACM transactions on intelligent systems and technology. 2012; 2:27, 1–27, 27.

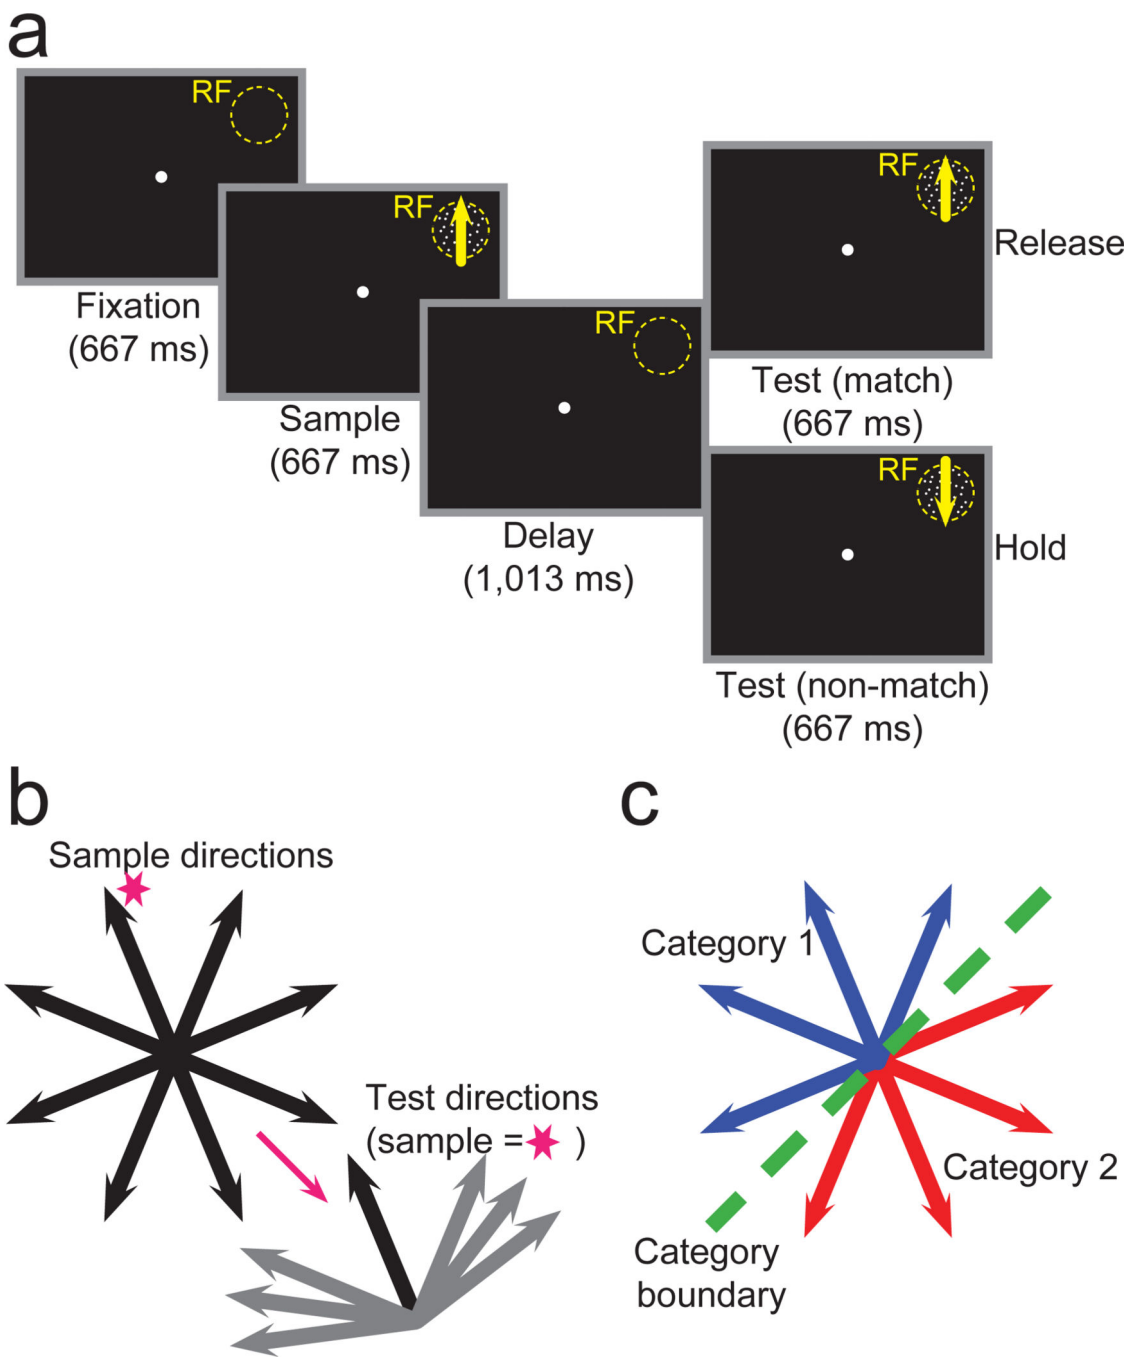

Figure 1. DMS and DMC tasks

(**a**) Monkeys performed a delayed matching task and indicated (by releasing a lever) whether sample and test stimuli, separated by a memory delay, were identical (DMS) or category (DMC) matches. If the sample and test stimuli weren't matches, the monkey was required to continue holding the lever throughout the test period and a second delay period until a second, matching test stimulus appeared. Stimuli were shown in a neuron's RF. (**b**) Monkeys viewed eight motion directions as sample stimuli during DMS task recordings (left). Test stimuli were either identical matches or 45°, 60°, or 75° from the sample stimulus. As an

example, the possible non-matching test directions are shown for the sample direction indicated by the magenta star. (**c**) Monkeys grouped the same eight motion directions used as sample directions during the DMS task into two categories (corresponding to the red and blue arrows) separated by a learned category boundary (dashed green line). Test and sample stimuli sets were identical.

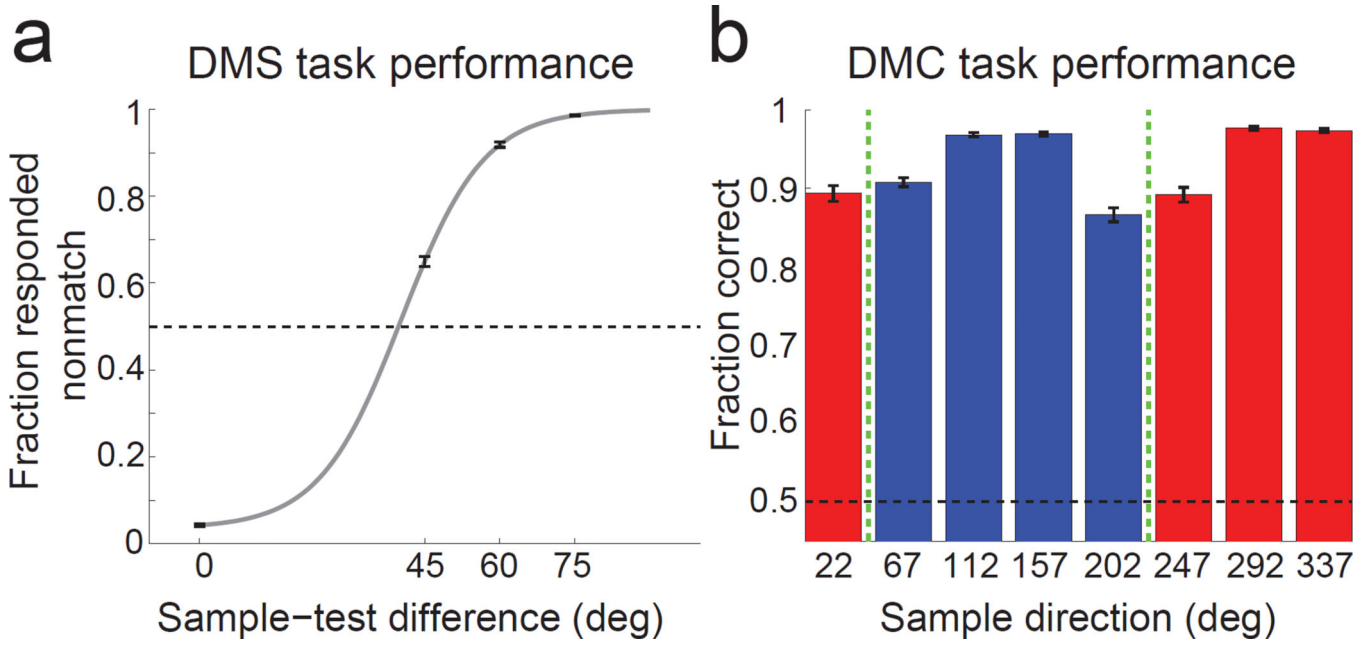

Figure 2. Behavioral performance

(**a**) The monkeys' average DMS performance (proportion reported as nonmatch) is shown as a function of sample-test difference. The gray line is a sigmoid curve fit to the behavioral data. The dashed black line indicates chance performance. Error bars indicate standard error of mean (SEM). (**b**) The monkeys' average DMC performance is shown as a function of sample direction. Colors correspond to the two categories and the green dashed lines correspond to the category boundary. Dashed black line indicates chance performance. Error bars indicate SEM.

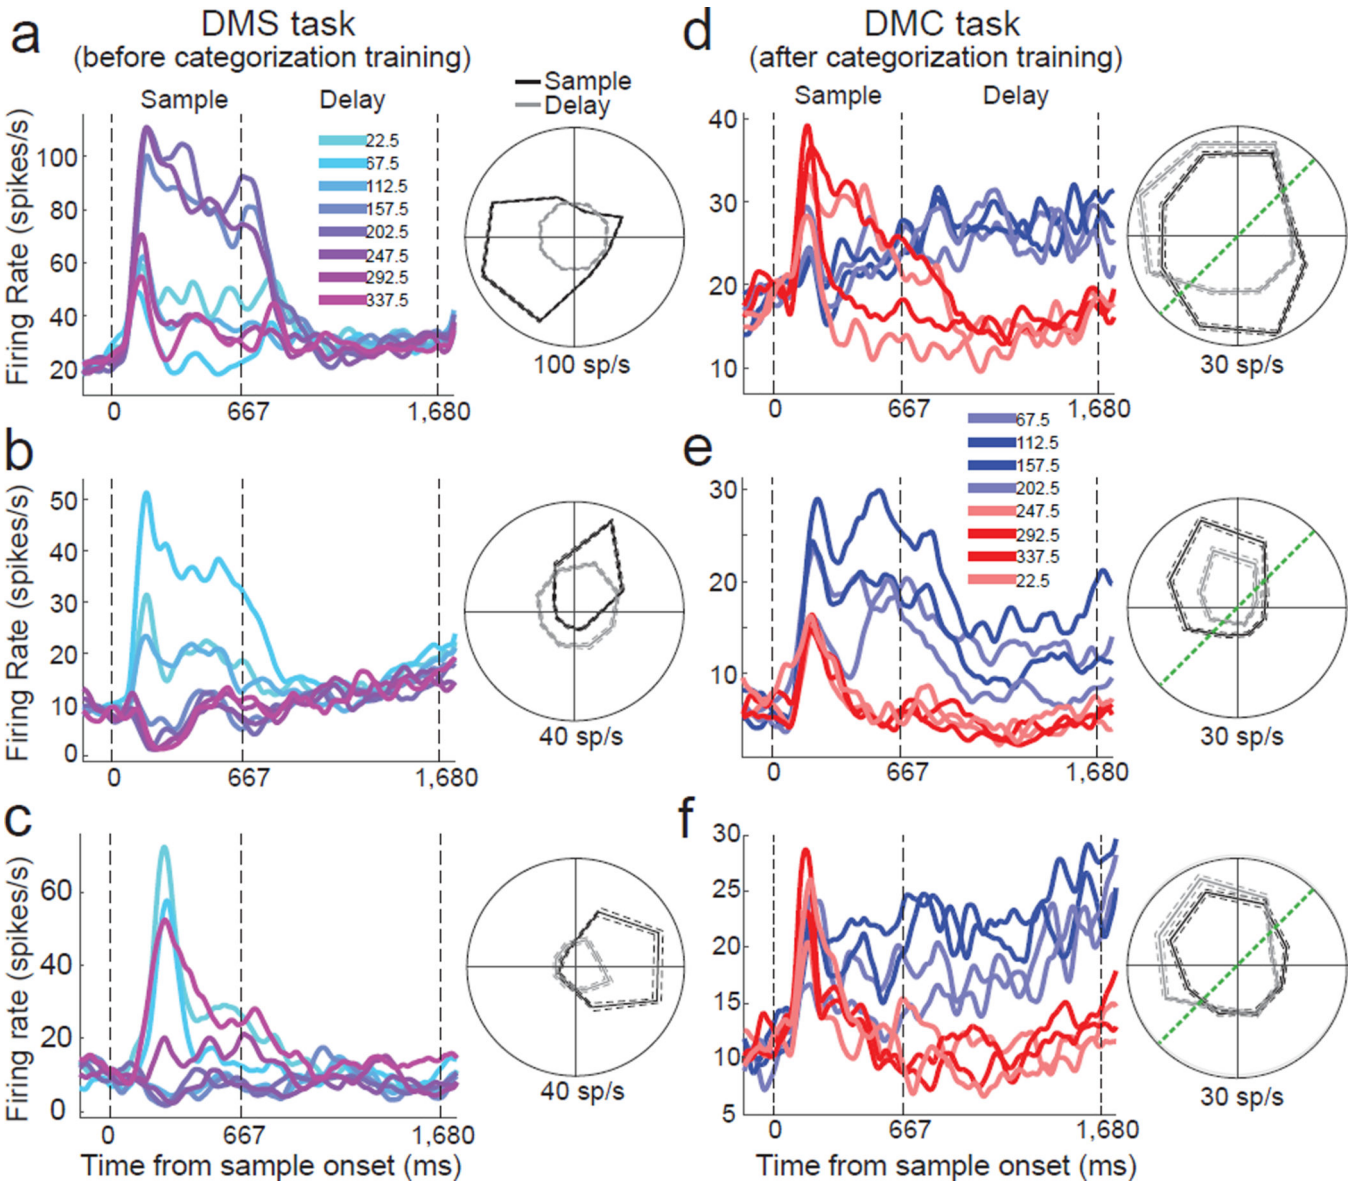

Figure 3. Example LIP neurons during the DMS and DMC tasks

Average activity evoked by 8 sample directions for three LIP neurons from the DMS (**a–c**) and DMC (**d–f**) tasks. Different traces indicate different sample directions and are colored according to their direction (and category membership for the DMC task). The three dashed, vertical lines represent the start of the sample epoch, the end of the sample epoch, and the end of the delay epoch, respectively. Polar plots are shown for average firing rates by sample direction during sample (black) and delay (gray) periods. In (**d–f**) dark blue and dark red traces indicate sample directions near the middle of categories 1 and 2, respectively, while light blue and light red traces indicate sample directions near the category boundary. All cells shown were direction selective during the sample period (one-way ANOVA, $P<0.01$). All cells except for (**a**) were also direction selective during the delay period. The dashed green lines on the polar plots indicate the position of the category boundary.

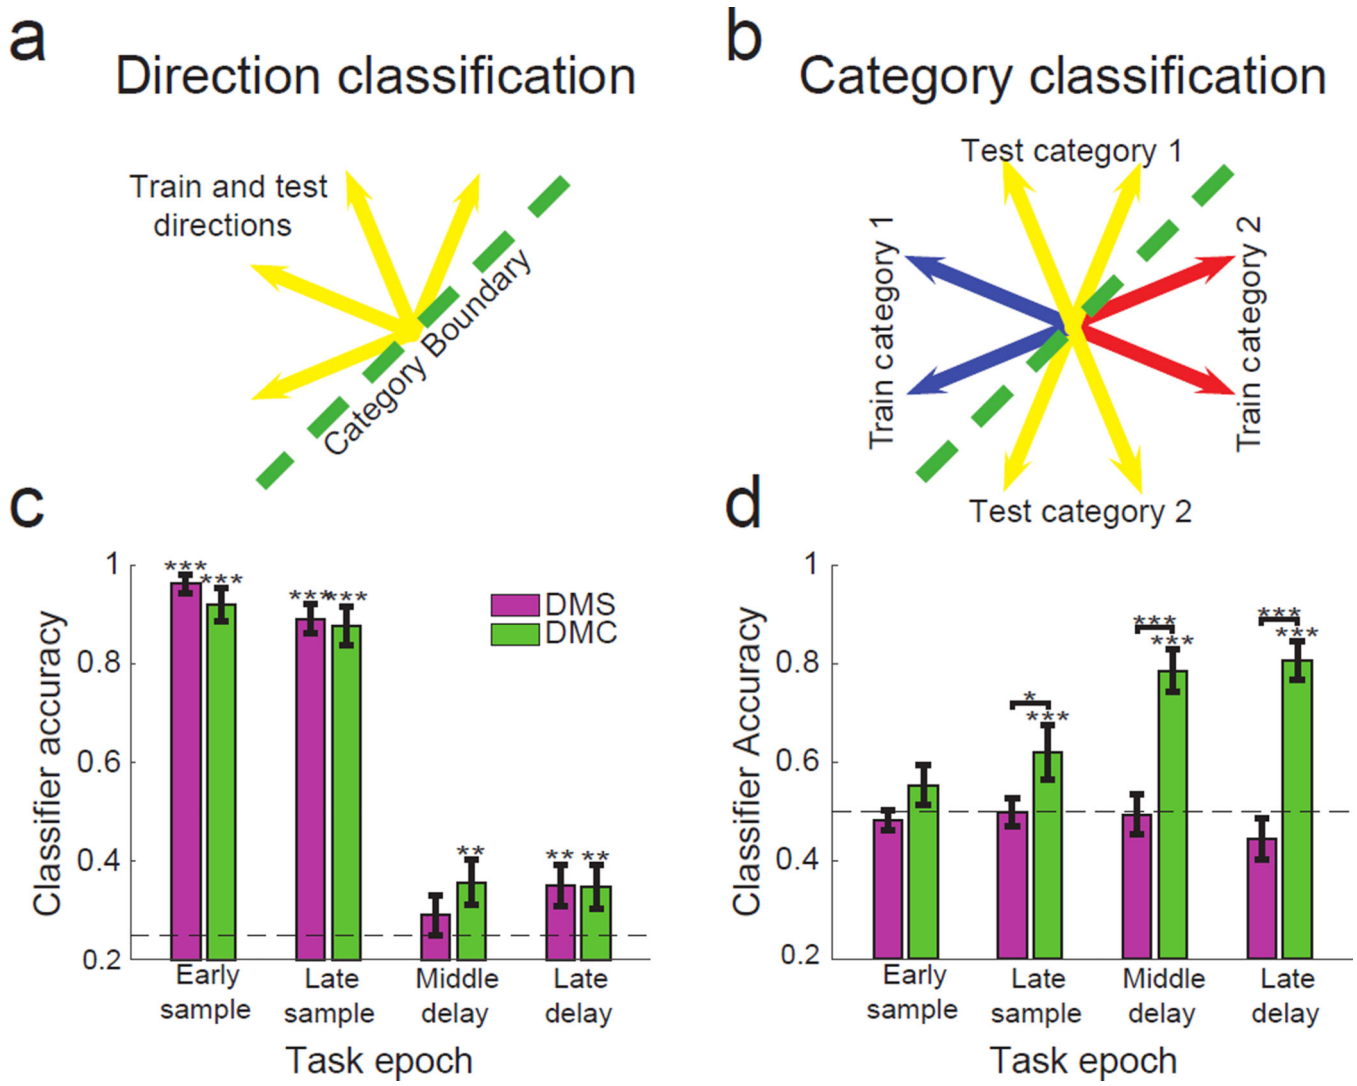

Figure 4. Population-level direction and category classification in LIP

(**a**) Diagram depicting how motion direction was classified independently of category. As an example, only neural responses to motion directions from category 1 (yellow arrows) are used to train the classifier. This classifier is then used to test which of the four motion directions (same yellow arrows) neural responses from these trials belong to. A second classifier was constructed using the four directions from the other category, and performance between these classifiers was averaged. (**b**) Diagram depicting how motion category was classified independently of direction. As an example, neural responses to two directions from category 1 (blue arrows) and two directions from category 2 (red arrows) are used to train the classifier. The classifier is then used to test which category neural responses to the four test directions (yellow arrows) belong to. A second classifier was constructed by switching the training and testing directions shown, and performance between these classifiers was averaged. (**c,d**) Stimulus selectivity during task epochs. (**c**) Performance of the direction classifier during the DMS (pink) and DMC (green) tasks. Chance accuracy is 0.25. (**d**) Performance of the category classifier during the DMS and DMC tasks. Chance

accuracy is 0.5. All epochs were 333 ms long. Error bars indicate SEM. *P<=0.05, **P<=0.01, ***P<=0.001, bootstrap.

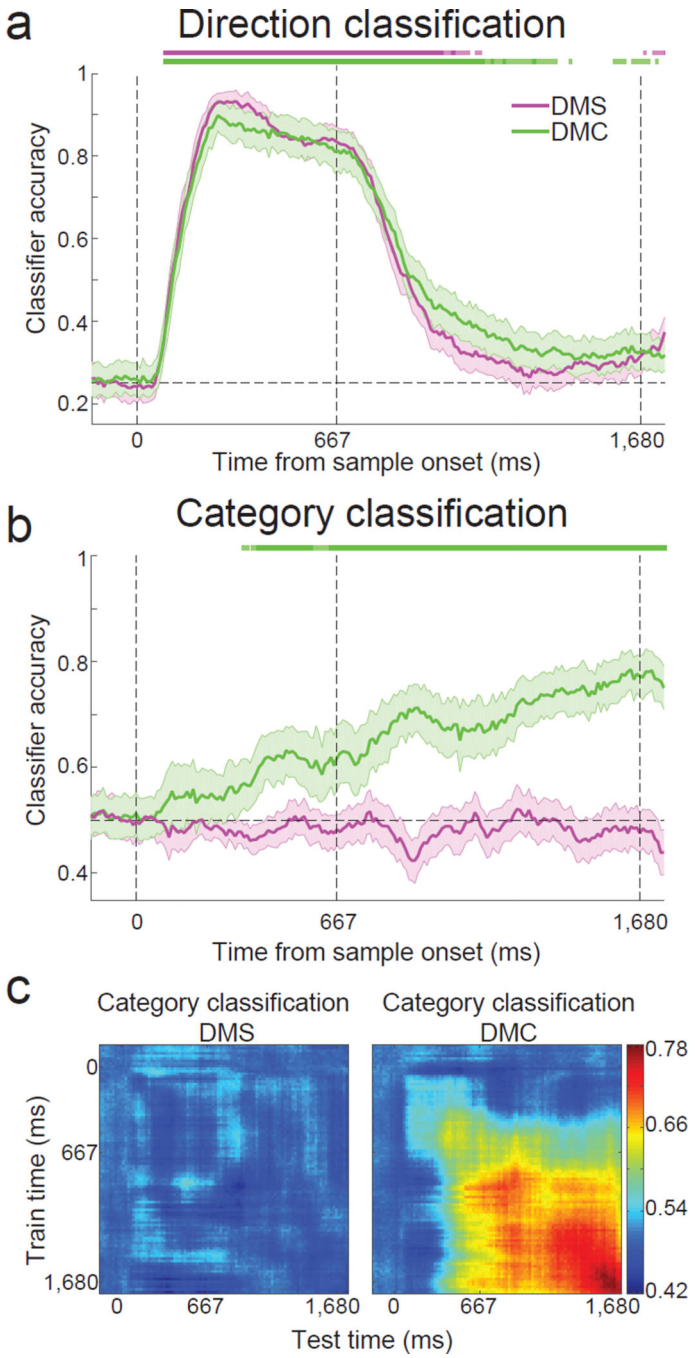

Figure 5. Time-course of direction and category classification in LIP

(**a,b**) The time course of direction and category selectivity in LIP during the two tasks was determined by computing classification accuracy as a function of time relative to sample onset using the (**a**) sample direction classifier and (**b**) sample category classifier. The three dashed, vertical lines represent the start of the sample epoch, the end of the sample epoch, and the end of the delay epoch, respectively. Error bars indicate SEM. The light and dark colored bars at the top of the figure indicate times at which classification accuracy during the DMS (pink) or DMC (green) task was significantly above chance (light, P<0.05, dark,

P<0.01, bootstrap). Sample direction classification accuracy was not significantly different between tasks at any point in the trial (P>0.05, bootstrap). (**c**) The stability of category encoding in LIP neurons was determined by training the classifier at one time point (y axis) and testing at a second time point (x axis). Classification accuracy is indicated by the color at each x–y coordinate. In the left panel, category classification accuracy during the DMS task remained at or near chance (0.5) for all training and testing time combinations. In the right panel, contiguous blocks of high category classification accuracy during the DMC task indicate temporally stable category decoding during the delay period—evident by the wide range of classifier training and testing times which yielded high category classification accuracy.

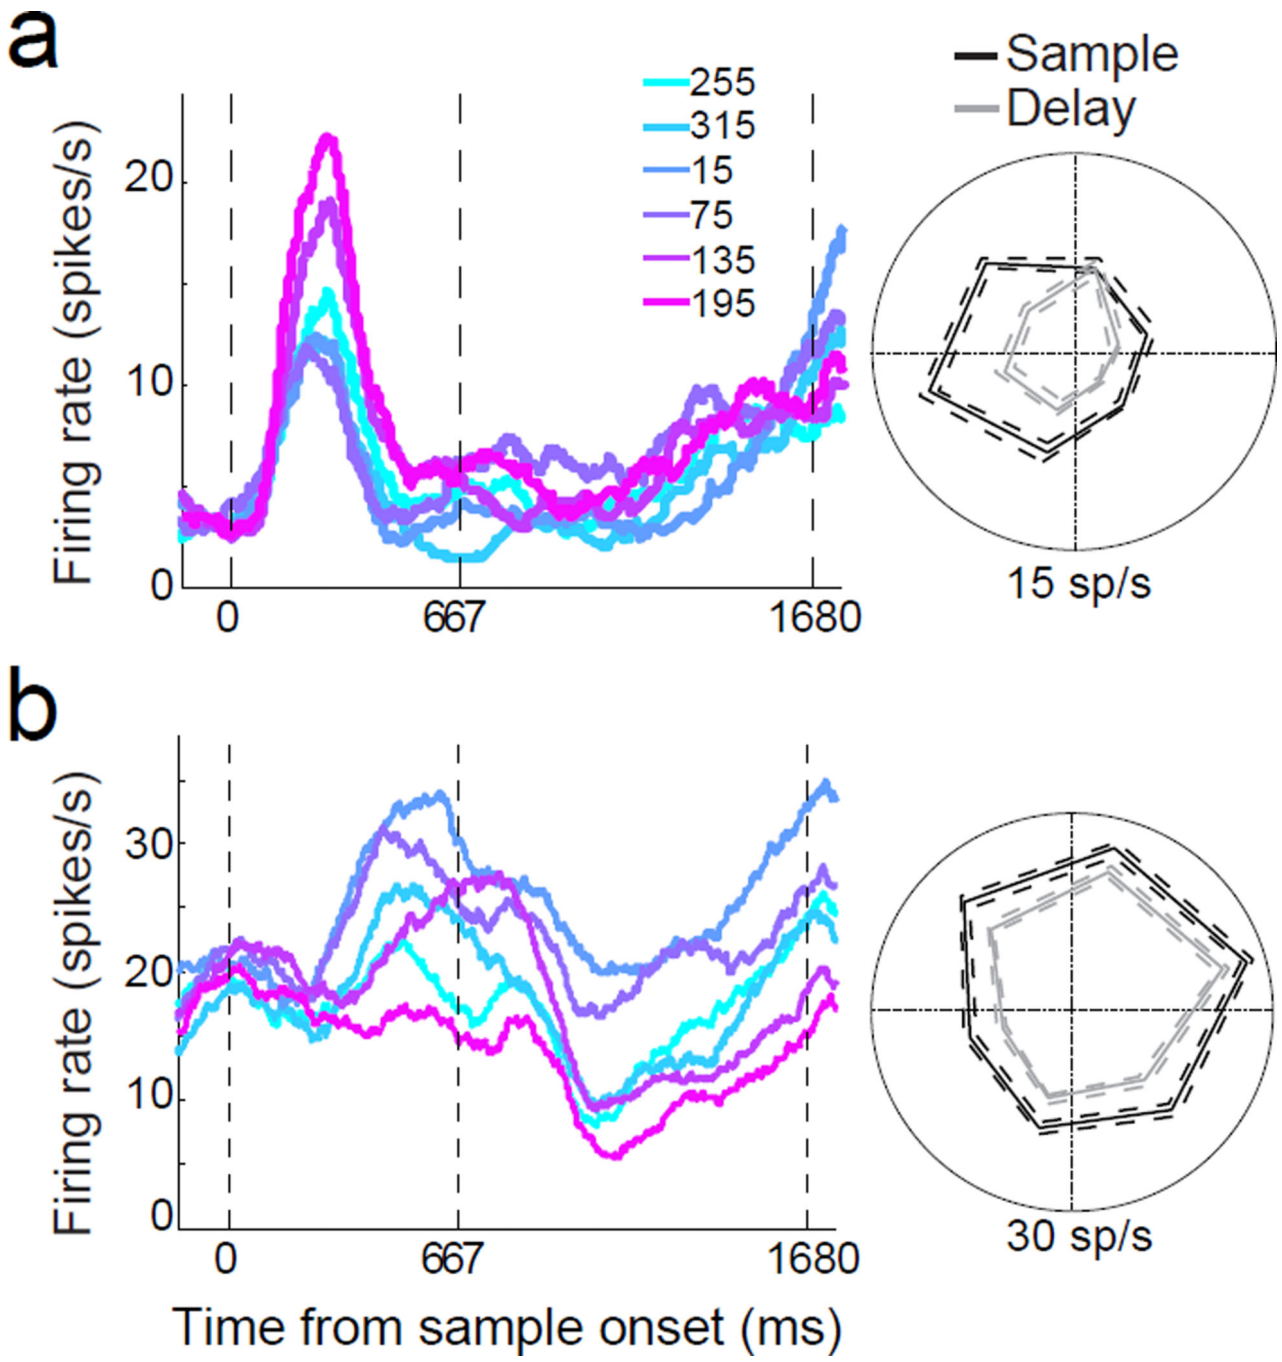

Figure 6. Example PPC and PFC neurons during the DMS task

Average activity evoked by 6 sample directions during the DMS task for a PPC neuron from monkey Q (**a**) and one PFC neuron from monkey W (**b**). Different traces indicate different sample directions and are colored according to their direction. The three dashed, vertical lines represent the start of the sample epoch, the end of the sample epoch, and the end of the delay epoch, respectively. Polar plots are shown for average firing rates by sample direction during sample (black) and delay (gray) periods. Dashed lines indicate SEM.

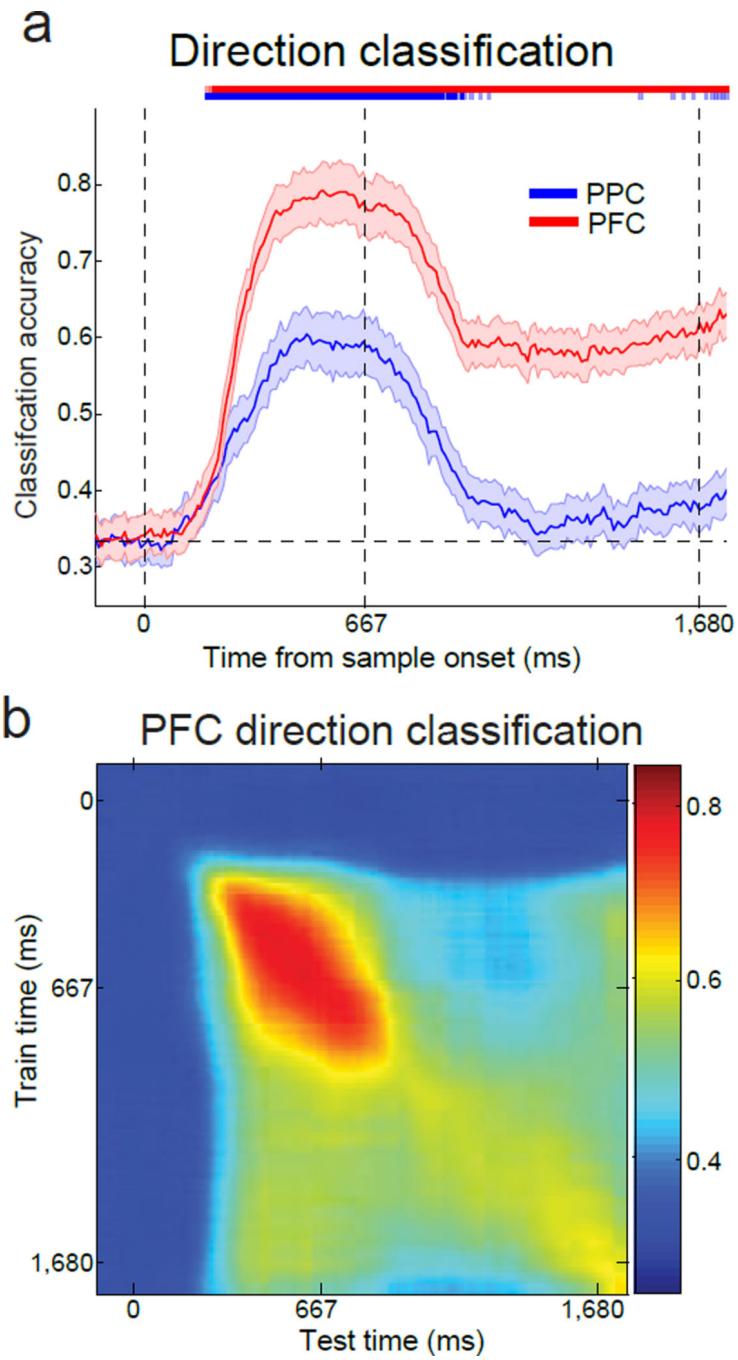

Figure 7. Time-course of direction classification in PPC and PFC

(**a**) Sample motion direction classification accuracy for Monkeys Q and W. Similar to Figure 5a, direction was decoded independent of category. The classification accuracy is shown for the PPC (blue curve) and the PFC (red curve). The three dashed, vertical lines represent the start of the sample epoch, the end of the sample epoch, and the end of the delay epoch, respectively. Error bars indicate SEM. The light and dark colored bars at the top of the panel indicate times at which classification accuracy for PPC (blue) and PFC (red) was significantly above chance (light, $P<0.05$, dark, $P<0.01$, bootstrap). (**b**) Similar to Figure 5c,

the stability of direction encoding in PFC was determined by training the classifier at one time point (y axis) and testing at a second time point (x axis). Classification accuracy is indicated by color at each x–y coordinate. Direction classification accuracy in PFC was above chance during both the sample and delay epochs, with stronger values near the diagonal, suggesting a combination of stable and dynamic direction encoding.