Local organization of spatial and shape information during working memory in the primate prefrontal cortex

Yunyi Sun[1], Wenhao Dang[2], Rye G. Jaffe[2], Christos Constantinidis[2,3,4]

1. Department of Biostatistics, Vanderbilt University Medical Center, Nashville TN 37203, USA.
2. Department of Biomedical Engineering, Vanderbilt University, Nashville TN 37235, USA.
3. Neuroscience Program, Vanderbilt University, Nashville TN 37235, USA.
4. Department of Ophthalmology and Visual Sciences, Vanderbilt University Medical Center, Nashville TN 37232, USA.

Abbreviated Title: PFC Organization

Address correspondence to:
Christos Constantinidis, PhD.
Christos.Constantinidis.1@vanderbilt.edu

The authors report no conflicts of interest during the research conducted.

ABSTRACT

While the current understanding of sensory and motor cortical areas has been defined topographical maps across the surface of these areas, higher cortical areas, such as the prefrontal cortex, seem to lack an equivalent organization, with only limited evidence of functional clustering of neurons with similar stimulus properties. We sought to examine whether neurons that represent similar spatial and object information are clustered in the monkey prefrontal cortex and whether such an organization only emerges as a result of training. We analyzed neurophysiological recordings from male macaque monkeys before and after they were trained to perform cognitive tasks. Neurons with similar spatial or shape selectivity were more likely than chance to be encountered at short distances from each other. This pattern of organization was present even in naïve animals, prior to any cognitive training. Our results reveal that prefrontal microstructure automatically supports orderly representations of spatial and object information.

INTRODUCTION

The prefrontal cortex (PFC) has long been recognized as the seat of higher cognitive functions, based on a plethora of evidence, including lesion, anatomical, and neurophysiological experiments in non-human primates (Goldman-Rakic 1988). A population of prefrontal neurons generate persistent activity selective for the properties of stimuli held in memory (Constantinidis et al. 2018). Computational models typically simulate such persistent activity by networks of units with recurrent connections between units with similar tuning for spatial location; this organizations allows activity initiated by a stimulus to be maintained in the network during working memory (Wimmer et al. 2014; Hart and Huk 2020).

The prefrontal cortex lacks an obvious retinotopic organization (Kiani et al. 2015; Leavitt et al. 2017), or other type of systematic topography of sensory information that would be analogous to the orderly arrangement of units in these artificial networks – or in the sensory and motor cortices. Only some local organization has been previously described in the prefrontal cortex, for spatial information, with nearby neurons representing the same or adjacent spatial locations (Constantinidis et al. 2001; Arnsten 2013). It is unclear therefore how neurons with similar selectivity could organize into functional circuits.

Newer theories of working memory, motivated primarily by results of human imaging studies have proposed that working memory information is represented in the sensory cortex, instead (Harrison and Tong 2009; Albers et al. 2013; Ester et al. 2013; Xing et al. 2013; Sreenivasan et al. 2014). According to this viewpoint, prefrontal cortex

plays a control or supervisory role, perhaps highlighting the locations of remembered stimuli, rather than maintaining their identity in working memory itself (Curtis and D'Esposito 2003; Tong and Pratte 2012; D'Esposito and Postle 2015). It is possible, therefore, that any local organization within the prefrontal cortex is limited to spatial information, alone.

We were thus motivated to examine the local organization of prefrontal neurons for the maintenance of spatial and object information in memory. We first replicated the analysis of spatial information of prior studies (Constantinidis *et al*. 2001), from a different set of experimental data. We then applied the same analysis to object information. Additionally, we examined how this information is altered by training to perform a working memory task, by analyzing data recorded from monkeys naïve to training, and after they were trained to perform spatial and object working memory tasks.

METHODS

Animals: Data obtained from six male rhesus monkeys (*Macaca mulatta*), ages 5–9 years old, weighing 5–12 kg, as previously documented (Riley et al. 2018), were analyzed in this study. None of these animals had any prior experimentation experience at the onset of our study. Monkeys were either single-housed or pair-housed in communal rooms with sensory interactions with other monkeys. All experimental procedures followed guidelines set by the U.S. Public Health Service Policy on Humane Care and Use of Laboratory Animals and the National Research Council's Guide for the Care and Use of Laboratory Animals and were reviewed and approved by the Wake Forest University Institutional Animal Care and Use Committee.

Experimental setup: Monkeys sat with their heads fixed in a primate chair while viewing a monitor positioned 68 cm away from their eyes with dim ambient illumination and were required to fixate on a 0.2° white square appearing in the center of the screen. In order to receive a liquid reward (typically fruit juice), the animals maintained fixation on the square while visual stimuli were presented either at a peripheral location or over the fovea. Any break of fixation immediately terminated the trial and no reward was given. Eye position was sampled at 240 Hz, monitored on line, and recorded with an infrared eye position scanning system (model RK-716; ISCAN, Burlington, MA). The visual stimulus display, monitoring of eye position, and synchronization of stimuli with neurophysiological data were performed with in-house software (Meyer and

Constantinidis 2005) implemented on the MATLAB environment (Mathworks, Natick, MA), utilizing the Psychophysics Toolbox (Brainard 1997).

Pre-training task: Following a brief period of fixation training and acclimation to the stimuli, monkeys were required to fixate on a center position while stimuli were displayed on the screen. The stimuli shown in the pre-training passive spatial task were white 2° squares, presented in one of nine possible locations arranged in a 3 × 3 grid with 10° distance between adjacent stimuli. The stimuli shown in the pre-training passive feature task were white 2° geometric shapes drawn from a set comprising a circle, diamond, the letter H, the hashtag symbol, the plus sign, a square, a triangle, and an inverted Y-letter. These stimuli could also be presented in one of nine possible locations arranged in a 3 × 3 grid with 10° distance between adjacent stimuli.

Presentation began with a fixation interval of 1 s where only the fixation point was displayed, followed by 500 ms of stimulus presentation (referred to hereafter as cue), followed by a 1.5 s “delay” interval (referred to hereafter as delay1) where, again, only the fixation point was displayed. A second stimulus (referred to hereafter as sample) was subsequently shown for 500 ms. For the spatial stimulus set, this second stimulus would be either identical in location to the initial stimulus, or diametrically opposite the first stimulus. For the feature set, this second stimulus would appear in the same location to the initial stimulus and would either be an identical shape or the corresponding non-match shape (each shape was paired with one non-match shape). Only one nonmatch stimulus was paired with each cue, so that the number of match and nonmatch trials were balanced in each set. In both the spatial and feature task, this second stimulus display was

followed by another “delay” period (delay2) of 1.5 s where only the fixation point was displayed. Successful holding of fixation until the end of the second delay period resulted in reward delivery. The location and identity of stimuli was of no behavioral relevance to the monkeys during the “pre-training” phase, as fixation was the only necessary action for obtaining reward.

Post-training task: Four of the six monkeys were trained to complete spatial and feature working memory tasks. These tasks involved presentation of identical stimuli as the spatial and feature tasks during the “pre-training” phase, but now monkeys were required to remember the spatial location and/or shape of the first presented stimulus, and report whether the second stimulus was identical to the first or not, via saccading to one of two target stimuli (green for matching stimuli, blue for non-matching). Each target stimulus could appear at one of two locations orthogonal to the cue/sample stimuli, pseudo-randomized in each trial.

Surgery and neurophysiology: A 20 mm diameter craniotomy was performed over the PFC and a recording cylinder was implanted over the site. The location of the cylinder was visualized through anatomical magnetic resonance imaging (MRI) and stereotaxic coordinates post-surgery. Electrode penetrations were aligned to the MRI image and mapped onto the cortical surface.

Neuronal recordings: Neural recordings were carried before and after training in the working memory tasks. The se of recordings analyzed here was previously used to

determine the properties of neurons before and after training in different PFC subdivisions and more details about the properties of these neurons can be found in previous reports (Meyer et al. 2011; Riley *et al.* 2018). Extracellular recordings were performed with multiple microelectrodes that were either glass- or epoxylite-coated tungsten, with a 250 μm diameter and 1–4 MΩ impedance at 1 kHz (Alpha-Omega Engineering, Nazareth, Israel). A Microdrive system (EPS drive, Alpha- Omega Engineering) advanced arrays of up to 8-microelectrodes, spaced 0.2–1.5 mm apart, through the dura and into the PFC. The signal from each electrode was amplified and band-pass filtered between 500 Hz and 8 kHz while being recorded with a modular data acquisition system (APM system, FHC, Bowdoin, ME). Waveforms that exceeded a user-defined threshold were sampled at 25 μs resolution, digitized, and stored for off-line analysis. Neurons were sampled in an unbiased fashion, collecting data from all units isolated from our electrodes, with no regard to the response properties of the isolated neurons. A semi-automated cluster analysis relied on the KlustaKwik algorithm, which applied principal component analysis of the waveforms to sort recorded spike waveforms into separate units (Harris et al. 2000). To ensure a stable firing rate in the analyzed recordings, we identified recordings in which a significant effect of trial sequence was evident at the baseline firing rate (ANOVA, $p < 0.05$), e.g., due to a neuron disappearing or appearing during a run, as we were collecting data from multiple electrodes. Data from these sessions were truncated so that analysis was only performed on a range of trials with stable firing rate. Less than 10% of neurons were corrected in this way. Identical data collection procedures, recording equipment, and spike sorting algorithms were used before and after training in order to prevent any confounds.

Data analysis. Analysis was implemented with the MATLAB computational environment (R2020-2023, Mathworks, Natick, MA). To obtain tuning similarity analysis by optimal stimuli, we first identified neurons with significant tuning for the stimulus set being examined. A one-way ANOVA at the α=0.05 significance level was employed, separately for each epoch of the task (fixation, cue, delay1, sample, delay2 in Fig. 1). The preferred stimulus in each task epoch was then defined for the neuron as the stimulus (location/shape) that elicited the highest average firing rate in the corresponding epoch. We examined pairs of nearby neurons recorded simultaneously from separate microelectrodes, placed together at distances of 0.2 – 1.5 mm from each other. For each task epoch, we identified such pairs both neurons of which were selective for the stimuli, and determined if they had the same preferred (best) stimulus or not. The proportions of such neuron pairs among all tuned neuron pairs were determined in each epoch and we used a nonparametric test to test the statistical significance of this proportion. More specifically, we randomly paired cells in our pool that were not recorded simultaneously (from nearby locations), and then repeated the same analysis described above to determine a simulated proportion of cells with the same preference. This random pairing procedure was repeated 1000 times to get a full null distribution. We then counted the number of occurrences in the null distribution for more extreme values compared to the empirical value and divided that number by 1000 to get a probability measure, which would be the p-value of the empirical measurement.

In addition to determining the best stimulus of each pair, we also analyzed the tuning similarity of nearby neurons by quantifying the Euclidean distance defined by their average response to all stimuli used in each task. More specifically, for a given

neuron, a trial-average firing rate vector was created by averaging the trials recorded with each stimulus, resulting in a 9 by 1 vector for the spatial test, and an 8 by 1 vector for the feature set. Each neuron's firing rate vector for a certain task epoch was then normalized by the highest value in the vector, to control for firing rate difference between neurons. For a given nearby neuron pair, and for a given task epoch, we calculated the Euclidean distance in the 9- (or 8-) dimensional space using the built-in MATLAB function "pdist2". If both cells in the pair displayed selectivity in more than one task epoch, the averaged tuning distance for all task epochs in which significant selectivity was present was used. We used a nonparametric test for testing the statistical significance of the Euclidean distance measurement. We again randomly paired neurons from different recordings, without having been recorded simultaneously. For a given neuron in this cell pool, we randomly chose one of the four epochs to get data from. This is the first source of our randomness introduced to the bootstrap analysis. Then we constructed a random cell pair pool containing all possible combinations of cells in the pool. And lastly we randomly draw the same number of pairs of cells as in the empirical pair dataset to calculate one value of the mean Euclidean distance for the null distribution. We had 1000 such iterations to generate the full null distribution.

To visualize the representational similarity of different stimuli at the population level, we used the MATLAB function "linkage" and "dendrogram". Specifically, we created an nine by N matrix (or eight by N matrix for the shape stimuli), where N is the number of unique neurons in the empirical neighboring-neuron pool as described above. Each element in the matrix represents the average neural response to 9 (or 8) stimuli. We then use the MATLAB function "pdist2" to create an 9 x 9 (or 8 by 8) pairwise distance

matrix representing the pairwise distance between the neuronal representation of the stimuli. Lastly, the distance matrix served as the input to the MATLAB “linkage” function to create a hierarchical cluster tree.

RESULTS

Database

Extracellular neurophysiological recordings were collected from the lateral PFC of six male monkeys before and after they were trained to perform the match/nonmatch working memory tasks (Meyer *et al.* 2011; Riley *et al.* 2018). The task required them to view two stimuli appearing in sequence with an intervening delay period between them, and after a second delay period to report whether or not the second stimulus was identical to the first (Fig. 1). The two stimuli could differ in terms of their location (spatial task, Fig. 1A-B) or shape (feature task, Fig. 1C-D). If the second stimulus was identical to the first, monkeys were required to saccade towards a green choice target that appeared at the end of the trial. Otherwise, they had to saccade to a blue target. The locations of the choice targets varied randomly from trial to trial and the monkey could not plan a motor response prior to their appearance.

We collected neurophysiological data before and after the monkeys learned to perform these tasks. A total of 1617 cells from six monkeys were recorded while the animals viewed the spatial stimuli passively prior to training, and 1493 cells from five monkeys were recorded while the animals viewed the feature stimuli passively, prior to training to perform the task (“pre-training” phase). A total of 1104 cells from three of the same monkeys were collected after the monkeys were trained and were performing the spatial working memory task. Another 1091 cells from two of the same monkeys were collected after the monkeys were trained and performed the feature working memory tasks (“post-training” phase).

Organization of Spatial Selectivity

We first identified neurons tuned for the spatial location of stimuli, evidenced by significant selectivity for the nine spatial stimuli (1-way ANOVA, $p<0.05$). Percentages of neurons that exhibited such significant selectivity at different time epochs of the task (plus the baseline, fixation period, as a control) are shown in Fig. 2. Spatial selectivity generally increased after training, as we have reported previously (Riley *et al.* 2018), with approximately 15% of all neurons recorded exhibiting selectivity during any task period.

We relied on this dataset to examine how neurons selective for spatial stimuli are arranged in local prefrontal circuits. We thus examined pairs of neurons recorded simultaneously from different microelectrodes, spaced 0.2 – 1.5 mm apart from each (mean distance = 0.77 mm). For 499 such pairs (222 recorded before and 277 recorded after training), both neurons exhibited significant selectivity in the same task epoch. We could thus test that neurons that had the same spatial preference (best response for the same location) were encountered more often than would be expected by chance (Fig.2C), which was evaluated using a bootstrap task, randomly pairing neurons recorded at different sessions and random positions. We indeed determined that neurons with the same preference were significantly higher than expected at multiple task epochs: cue, and sample period, prior to training ($p<0.001$ in each case), and cue and delay period after training ($p=0.012$ and $p<0.001$, respectively).

This analysis considers only the best location, however, the entire tuning function of a neuron provides a more complete picture of neuronal tuning. We thus refined our analysis to consider the vector defined by firing rates for all eight locations for each neuron. We then computed the Euclidean distance between the vectors of two neurons,

examining again only task epochs in which both neurons exhibited significant spatial selectivity. For pairs with significant selectivity in more than one task epoch, we computed distances in each epoch and then averaged them to obtain a single distance value for each neuron. To ensure that neurons with higher firing rates did not appear to exhibit longer distances, we normalized firing rates before computing distances. The mean distance computed in this manner was 1.02 prior to training and 1.11 after training (Fig. 3). Both values were significantly lower than would be expected by chance (permutation test, $p<0.001$ in each case).

The pattern of responses to different stimuli also allowed us to determine how different spatial locations are represented across neurons. A hierarchical cluster analysis mostly captured the topology of stimulus location, with adjacent locations being represented by nearby clusters (Fig. 4). Prior to training, the cluster corresponding to the shorter distance in stimulus activity involved locations 4 and 5, which represent adjacent stimuli on the screen (Fig. 1). Successive, higher-order clusters involved stimuli 6, 3, 9, 7, 1, 8, and 2, each one of which were adjacent with the preceding cluster of locations. A similar ordering of stimulus locations was present after training, although the absolute order of cluster distances changed somewhat.

Organization of Feature Selectivity

These results of the spatial analysis mostly replicated earlier studies which suggested a local organization of spatial information across the surface of the prefrontal cortex, using different spatial working memory tasks (Constantinidis *et al.* 2001). Since it has been argued that spatial representations are fundamentally different than object working

memory (Curtis and D'Esposito 2003; Tong and Pratte 2012; D'Esposito and Postle 2015), we were interested to test whether neuronal selectivity for shapes exhibits some organization or not.

We thus repeated our analysis of neuronal selectivity for shape information. The percentage of neurons with significant selectivity for the eight geometric shapes of our dataset evaluated with a 1-way ANOVA test ($p<0.05$) is shown in Fig. 5. Although shape selectivity was lower compared to spatial selectivity, neurons selective for shape were found both before and after training, with approximately 10% of all neurons recorded exhibiting shape selectivity at any task epoch.

A total of 242 pairs of neurons were recorded simultaneously from separate electrodes spaced 0.2 – 1.5 mm from each other, both neurons of which exhibited shape selectivity in the same task epoch (Fig.5C). We again examined whether these neurons were more likely to exhibit the same preferred stimulus than would be expected by chance. For this dataset, a higher than chance same-shape preference was observed in the cue period, after training (permutation test, $p=0.043$).

We compared the overall similarity of shape tuning between simultaneously recorded neurons by computing their normalized Euclidean distance, as we did for the spatial stimuli. The mean Euclidean distance was 0.86 prior to training and 0.88 after training. These distances were significantly smaller than would be expected by chance (permutation test, $p<0.001$, in each case). These results suggest that shape selectivity too, is not represented randomly across the surface of the cortex. Prefrontal neurons tuned for stimulus shape are present prior to any training as well as after training and organized in patches of cortex with similar shape selectivity.

Finally, we determined how different shapes are represented across the population of prefrontal neurons (Fig. 7). There is no obvious way to group the different shapes in analogy to spatial locations. Clustering was generally similar before and after training, although the absolute order of cluster distances changed somewhat after training.

DISCUSSION

Our analysis revealed a local organization of both spatial and non-spatial information across the surface of the prefrontal cortex, with neurons 1.5 mm of each other representing significantly more similar stimuli. These results are consistent with previous studies that have suggested a local organization of spatial representations across the prefrontal cortex, in monkeys trained to perform spatial working memory tasks (Constantinidis *et al.* 2001; Leavitt *et al.* 2017). Moreover, we now reveal a similar pattern of local organization for object representation, as well. Furthermore, we now document that this organization is largely present even prior to any training to perform a working memory task that required these stimuli to be recognized and maintained in memory.

Organization of information across the prefrontal cortex

The prefrontal cortex is critical for higher cognitive functions, including the maintenance of information in working memory (D'Esposito and Postle 2015; Riley and Constantinidis 2016). Specifically, this information is thought to be maintained through persistent discharges in a network of brain areas centered on the prefrontal cortex (Constantinidis *et al.* 2018). Computational models suggest that persistent activity is sustained in neural circuits by virtue of re-entrant connections between neurons with similar tuning for stimulus properties (Compte et al. 2000; Wimmer *et al.* 2014). Neural activity generated by a stimulus presentation is maintained in the network, which behaves as a continuous attractor (Seung et al. 2000). Until now however, the specific organization of these

neuronal networks across the surface of the prefrontal cortex has remained unresolved. No obvious, retinotopic organization of visual space has been described across the surface of the prefrontal cortex, or other higher-order cortical areas implicated in the maintenance of working memory, such as the posterior parietal cortex (Qi et al. 2015). Some indications exist for at least some local PFC organization. For example, electrode penetrations that descent into the principal sulcus, parallel to its surface, suggest an orderly progression of receptive field locations (Arnsten 2013). Our current results suggest that selectivity for stimulus properties is an organizing principle in the neural circuits of the prefrontal cortex, as speculated by early models of the cellular basis of working memory (Goldman-Rakic 1995).

Other findings also suggest at least some coarse localization of sensory information across domains. Whether object working memory is localized in a different subdivision of the prefrontal cortex than spatial working memory (ventral vs. dorsal) has been a matter of debate (Wilson et al. 1993; Rao et al. 1997), but at least a quantitative difference seems to be evident, with spatial information more prevalent in the dorsolateral prefrontal cortex than the ventrolateral prefrontal cortex (Meyer *et al*. 2011; Constantinidis and Qi 2018). Some organization for the time course of activation has also been revealed across cortical layers: neurons activated by visual stimuli are more likely to be encountered in intermediate layers; neurons with delay period activity in superficial; neurons with motor responses in deep, though this relative segregation is quantitative rather than qualitative (Markowitz et al. 2015; Zhu et al. 2023). Finally, an anterior-posterior gradient of stimulus properties has also been described so that posterior neurons are more selective for the sensory attributes of stimuli no matter what their context,

whereas more anterior neurons are much more likely to become activated only in the context of learned tasks (Riley et al. 2017; Riley *et al.* 2018).

Training effects

The activity of prefrontal neurons is highly plastic and responding to abstract variables of the task being executed (Fuster et al. 2000; Wallis et al. 2001; Mante et al. 2013). Neurophysiological studies of training have revealed that learning to perform cognitive tasks generally increases responsiveness to stimuli that are relevant for the task execution (Meyers et al. 2012; Constantinidis and Klingberg 2016). Prefrontal neurons, and particularly those in more anterior areas, become responsive to stimuli that previously failed to elicit a response (Qi et al. 2011; Riley *et al.* 2018). Recent human imaging results parallel these findings, confirming appearance of stimulus selectivity for learned stimuli after training (Miller et al. 2022). A variety of other changes have also been reported, including in variability and correlation of prefrontal discharges (Qi and Constantinidis 2012).

In view of these findings, it appeared possible that the local organization of the prefrontal cortex with respect to stimulus selectivity might only emerge after training. Contrary to this expectation, we found neurons with similar selectivity at nearby locations in the cortex, even in naïve monkeys, prior to any training. This result is consistent with previous findings suggesting that underlying network circuitry appears to connect neurons with similar properties, and that these functional modules of neurons remain fairly stable across different behavioral contexts or tasks (Kiani *et al.* 2015).

ACKNOWLEDGMENTS

Research reported in this paper was supported by the NIH National Eye Institute under award number R01 EY017077 and by the National Institute of Mental Health under award number R01 MH116675. We wish to acknowledge Chrissy Suell for technical help; Junda Zhu, and Zhengyang Wang for helpful comments on the manuscript.

REFERENCES

Albers AM, Kok P, Toni I, Dijkerman HC, de Lange FP. 2013. Shared representations for working memory and mental imagery in early visual cortex. Current biology : CB. 23:1427-1431.

Arnsten AF. 2013. The neurobiology of thought: the groundbreaking discoveries of Patricia Goldman-Rakic 1937-2003. Cereb Cortex. 23:2269-2281.

Brainard DH. 1997. The Psychophysics Toolbox. Spat Vis. 10:433-436.

Compte A, Brunel N, Goldman-Rakic PS, Wang XJ. 2000. Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. Cereb Cortex. 10:910-923.

Constantinidis C, Franowicz MN, Goldman-Rakic PS. 2001. Coding specificity in cortical microcircuits: a multiple electrode analysis of primate prefrontal cortex. J Neurosci. 21:3646-3655.

Constantinidis C, Funahashi S, Lee D, Murray JD, Qi XL, Wang M, Arnsten AFT. 2018. Persistent Spiking Activity Underlies Working Memory. J Neurosci. 38:7020-7028.

Constantinidis C, Klingberg T. 2016. The neuroscience of working memory capacity and training. Nat Rev Neurosci. 17:438-449.

Constantinidis C, Qi XL. 2018. Representation of Spatial and Feature Information in the Monkey Dorsal and Ventral Prefrontal Cortex. Front Integr Neurosci. 12:31.

Curtis CE, D'Esposito M. 2003. Persistent activity in the prefrontal cortex during working memory. Trends in Cognitive Sciences. 7:415-423.

D'Esposito M, Postle BR. 2015. The cognitive neuroscience of working memory. Annu Rev Psychol. 66:115-142.

Ester EF, Anderson DE, Serences JT, Awh E. 2013. A Neural Measure of Precision in Visual Working Memory. Journal of Cognitive Neuroscience. 25:754-761.

Fuster JM, Bodner M, Kroger JK. 2000. Cross-modal and cross-temporal association in neurons of frontal cortex. Nature. 405:347-351.

Goldman-Rakic PS. 1988. Topography of cognition: parallel distributed networks in primate association cortex. Annu Rev Neurosci. 11:137-156.

Goldman-Rakic PS. 1995. Cellular basis of working memory. Neuron. 14:477-485.

Harris KD, Henze DA, Csicsvari J, Hirase H, Buzsaki G. 2000. Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. J Neurophysiol. 84:401-414.

Harrison SA, Tong F. 2009. Decoding reveals the contents of visual working memory in early visual areas. Nature. 458:632-635.

Hart E, Huk AC. 2020. Recurrent circuit dynamics underlie persistent activity in the macaque frontoparietal network. Elife. 9.

Kiani R, Cueva CJ, Reppas JB, Peixoto D, Ryu SI, Newsome WT. 2015. Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex. Neuron. 85:1359-1373.

Leavitt ML, Pieper F, Sachs AJ, Martinez-Trujillo JC. 2017. A Quadrantic Bias in Prefrontal Representation of Visual-Mnemonic Space. Cereb Cortex.1-17.

Mante V, Sussillo D, Shenoy KV, Newsome WT. 2013. Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature. 503:78-84.

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.

Meyer T, Constantinidis C. 2005. A software solution for the control of visual behavioral experimentation. J Neurosci Methods. 142:27-34.

Meyer T, Qi XL, Stanford TR, Constantinidis C. 2011. Stimulus selectivity in dorsal and ventral prefrontal cortex after training in working memory tasks. J Neurosci. 31:6266-6276.

Meyers EM, Qi XL, Constantinidis C. 2012. Incorporation of new information into prefrontal cortical activity after learning working memory tasks. Proc Natl Acad Sci U S A. 109:4651-4656.

Miller JA, Tambini A, Kiyonaga A, D'Esposito M. 2022. Long-term learning transforms prefrontal cortex representations during working memory. Neuron. 110:3805-3819 e3806.

Qi XL, Constantinidis C. 2012. Correlated discharges in the primate prefrontal cortex before and after working memory training Eur J Neurosci. 36:3538-3548.

Qi XL, Elworthy AC, Lambert BC, Constantinidis C. 2015. Representation of remembered stimuli and task information in the monkey dorsolateral prefrontal and posterior parietal cortex. J Neurophysiol. 113:44-57.

Qi XL, Meyer T, Stanford TR, Constantinidis C. 2011. Changes in Prefrontal Neuronal Activity after Learning to Perform a Spatial Working Memory Task. Cereb Cortex. 21:2722-2732.

Rao SC, Rainer G, Miller EK. 1997. Integration of what and where in the primate prefrontal cortex. Science. 276:821-824.

Riley MR, Constantinidis C. 2016. Role of prefrontal persistent activity in working memory. Front Syst Neurosci. 9:181.

Riley MR, Qi XL, Constantinidis C. 2017. Functional specialization of areas along the anterior-posterior axis of the primate prefrontal cortex. Cereb Cortex. 27:3683–3697.

Riley MR, Qi XL, Zhou X, Constantinidis C. 2018. Anterior-posterior gradient of plasticity in primate prefrontal cortex. Nat Commun. 9:3790.

Seung HS, Lee DD, Reis BY, Tank DW. 2000. Stability of the memory of eye position in a recurrent network of conductance-based model neurons. Neuron. 26:259-271.

Sreenivasan KK, Vytlacil J, D'Esposito M. 2014. Distributed and Dynamic Storage of Working Memory Stimulus Information in Extrastriate Cortex. Journal of Cognitive Neuroscience. 26:1141-1153.

Tong F, Pratte MS. 2012. Decoding patterns of human brain activity. Annual review of psychology. 63:483-509.

Wallis JD, Anderson KC, Miller EK. 2001. Single neurons in prefrontal cortex encode abstract rules. Nature. 411:953-956.

Wilson FA, O Scalaidhe SP, Goldman-Rakic PS. 1993. Dissociation of object and spatial processing domains in primate prefrontal cortex. Science. 260:1955-1958.

Wimmer K, Nykamp DQ, Constantinidis C, Compte A. 2014. Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. Nat Neurosci. 17:431-439.

Xing Y, Ledgeway T, McGraw PV, Schluppeck D. 2013. Decoding working memory of stimulus contrast in early visual cortex. The Journal of neuroscience : the official journal of the Society for Neuroscience. 33:10301-10311.

Zhu J, Hammond BM, Constantinidis C. 2023. Laminar pattern of adolescent development changes in working memory neuronal activity. bioRxiv. doi.org/10.1101/2023.07.28.550982.

FIGURES

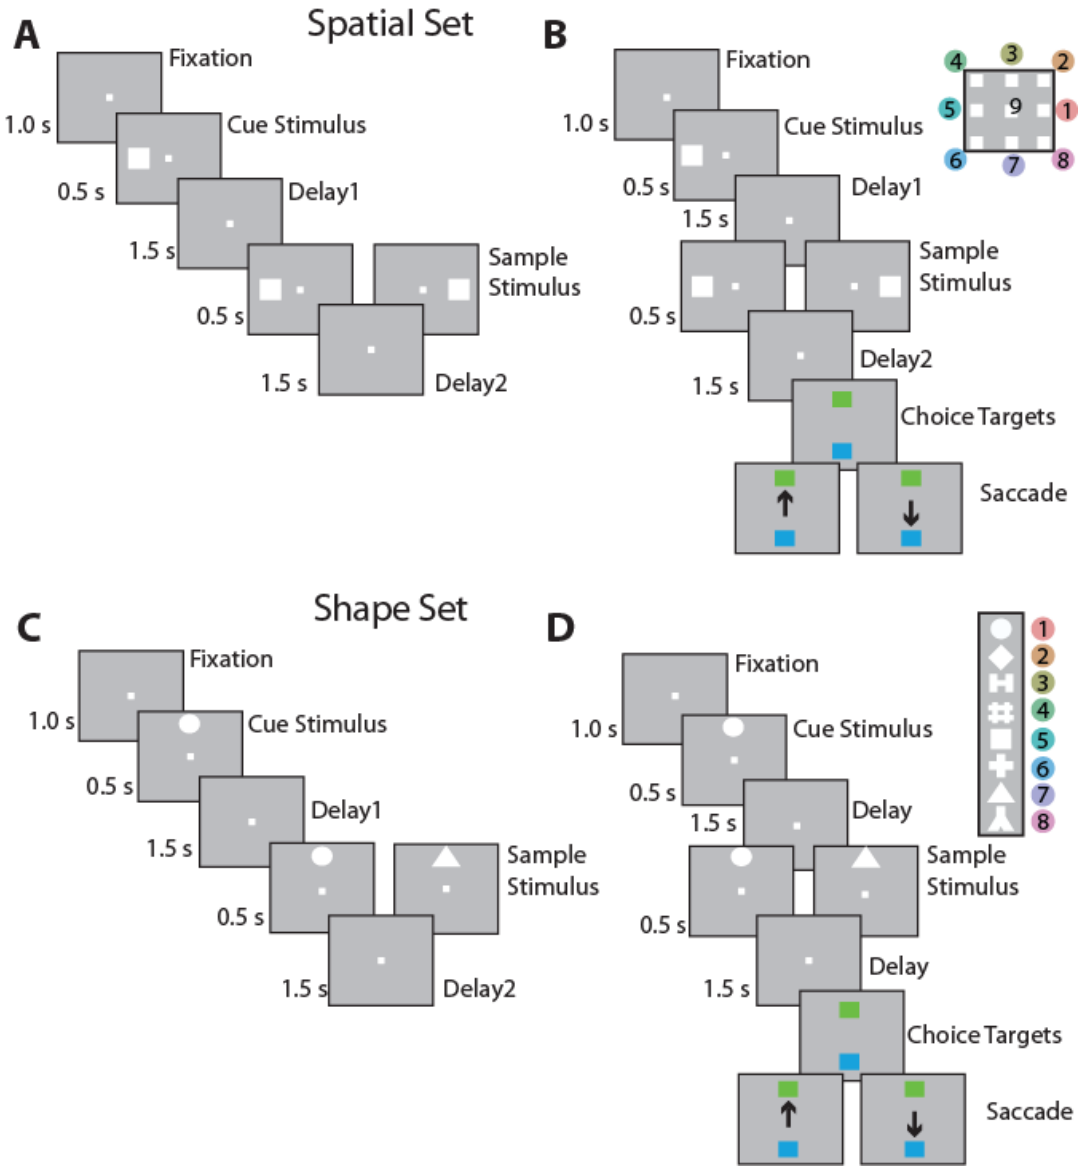

Figure 1. Task structure and stimuli used. A. Passive, spatial task. B. Active spatial task. nine possible cue locations in a session shown in the inset. The peripheral locations are numbered from 1-8, for reference in other figures. The animals were required to maintain center fixation throughout both active and passive task trials. At the end of active tasks trials however, monkeys were required to make a saccade to a green target if the stimuli matched or to a blue target if the stimuli did not match. C-D. Passive and active shape task. Eight possible shapes in a session are shown in the inset, numbered from 1-8, for reference in other figures. Stimuli in both tasks extended 2 degree of visual angle.

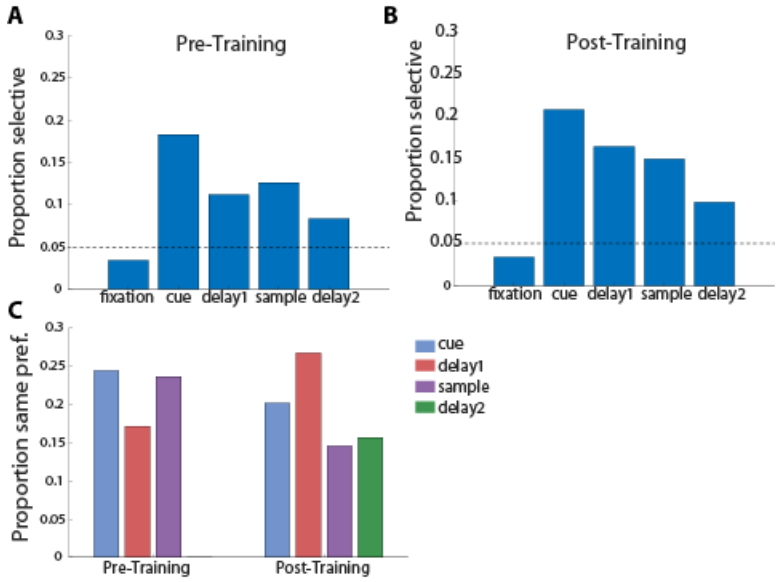

Figure 2. Proportion of neurons exhibiting significant spatial selectivity at different task epochs. A. Selective neurons prior to training. Sample size n=1661. B. Selective neurons post-training (n=1185 neurons). C. Proportion of neurons with the same spatial selectivity at different task intervals.

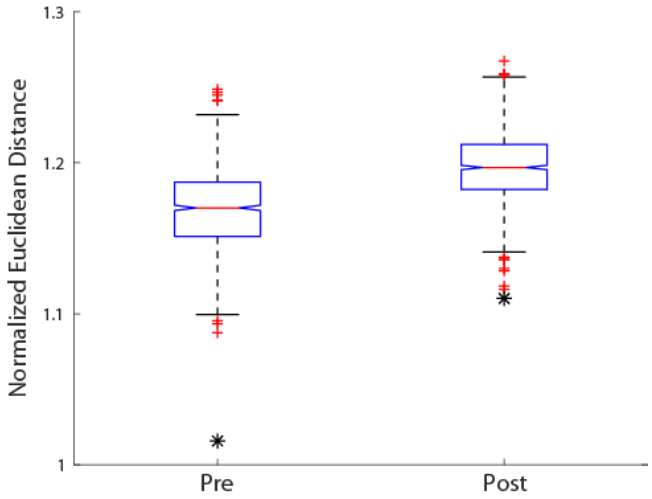

Figure 3. Normalized Euclidean distance of the representations of spatial stimuli among pairs of neighboring prefrontal neurons. The box borders represent the interquartile range (IQR) of the null distribution generated randomly. Whisker length indicates 1.5 interquartile range (IQR). Red crosses are outliers in the null distribution,. Black asterisks represent empirical mean Euclidean distance values. The representational distance between nearby neurons is smaller than the expected chance level.

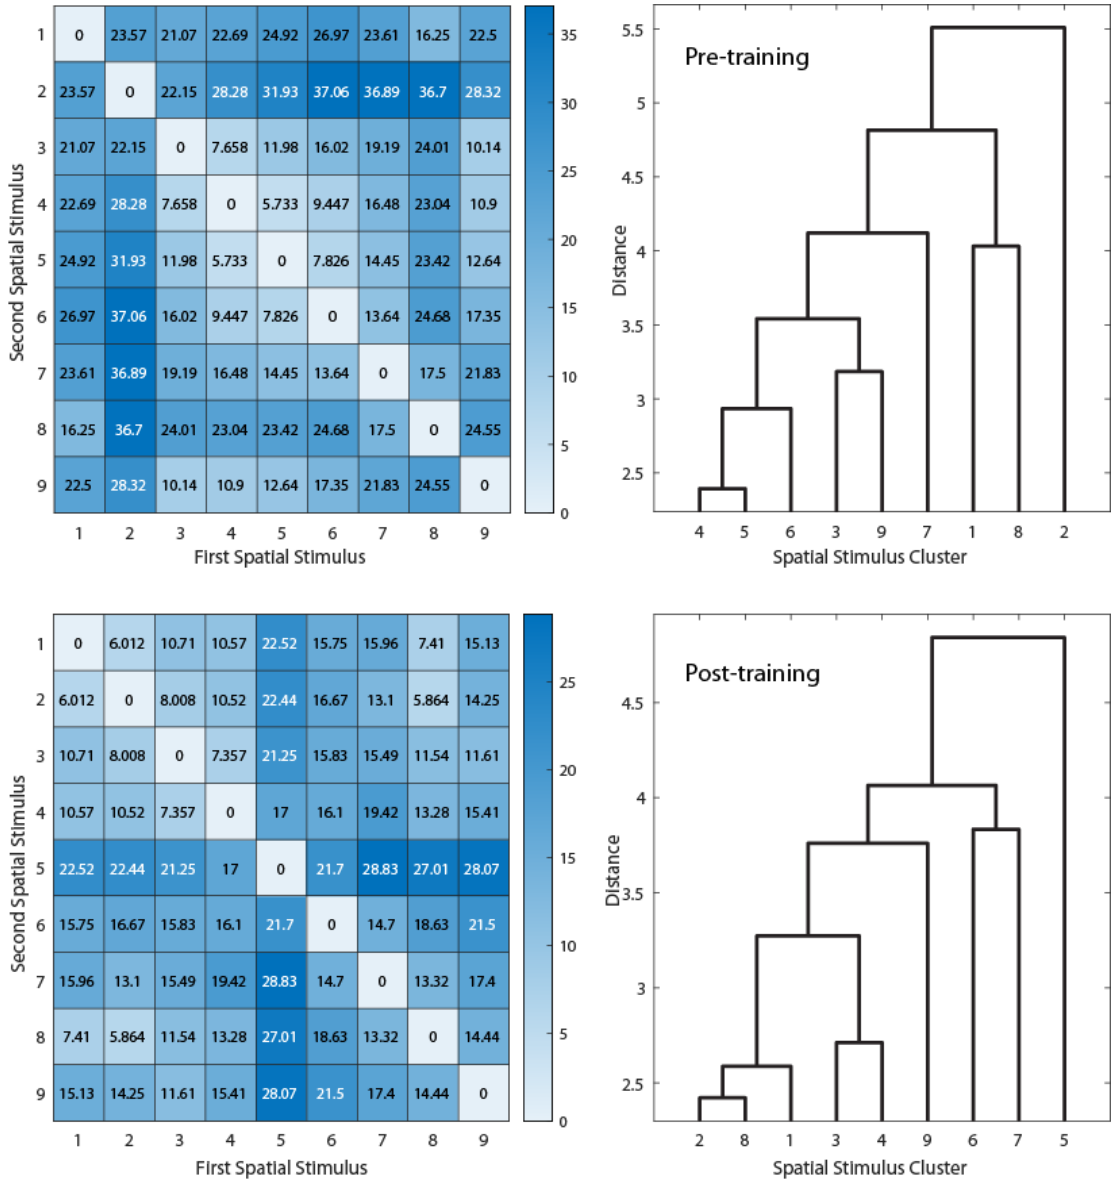

Figure 4. Heat maps and dendrogram representing the representational distance between any two stimuli used in the spatial task prior to training (top) and after training (bottom). Stimulus numbers refer to spatial locations as indicated in Fig. 1. Deeper color in the heat maps indicates less similarity between the neuronal responses elicited by two stimuli. Distance numbers are in arbitrary units.

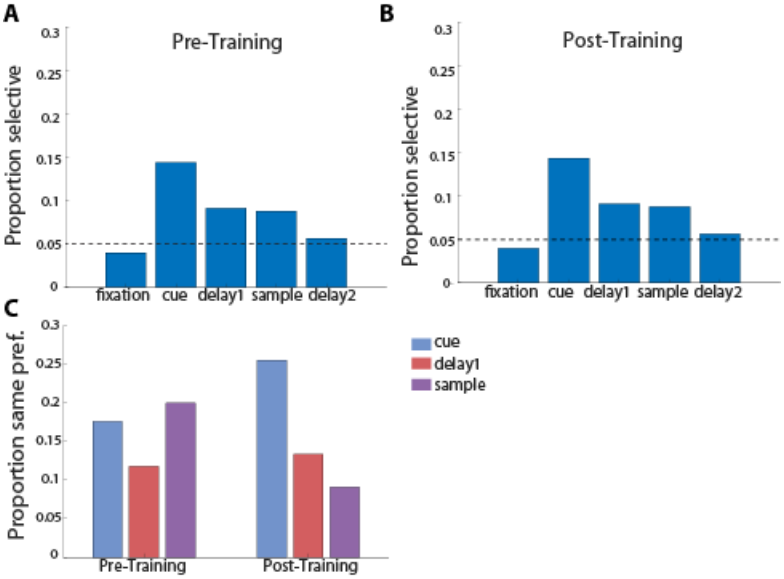

Figure 5. Proportion of neurons exhibiting significant shape selectivity at different task epochs. A. Selective neurons prior to training. Sample size n=1493. B. Selective neurons post-training (n=1091 neurons). C. Proportion of neurons with the same spatial selectivity at different task intervals.

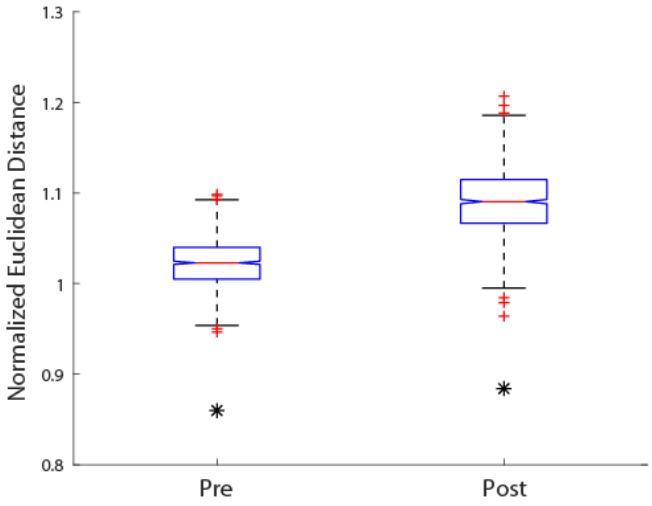

Figure 6. Normalized Euclidean distance of the representations of shape stimuli among pairs of neighboring prefrontal neurons. Conventions are the same as Fig. 3. In the Feature task, the representational distance between nearby neurons is smaller than the expected chance level.

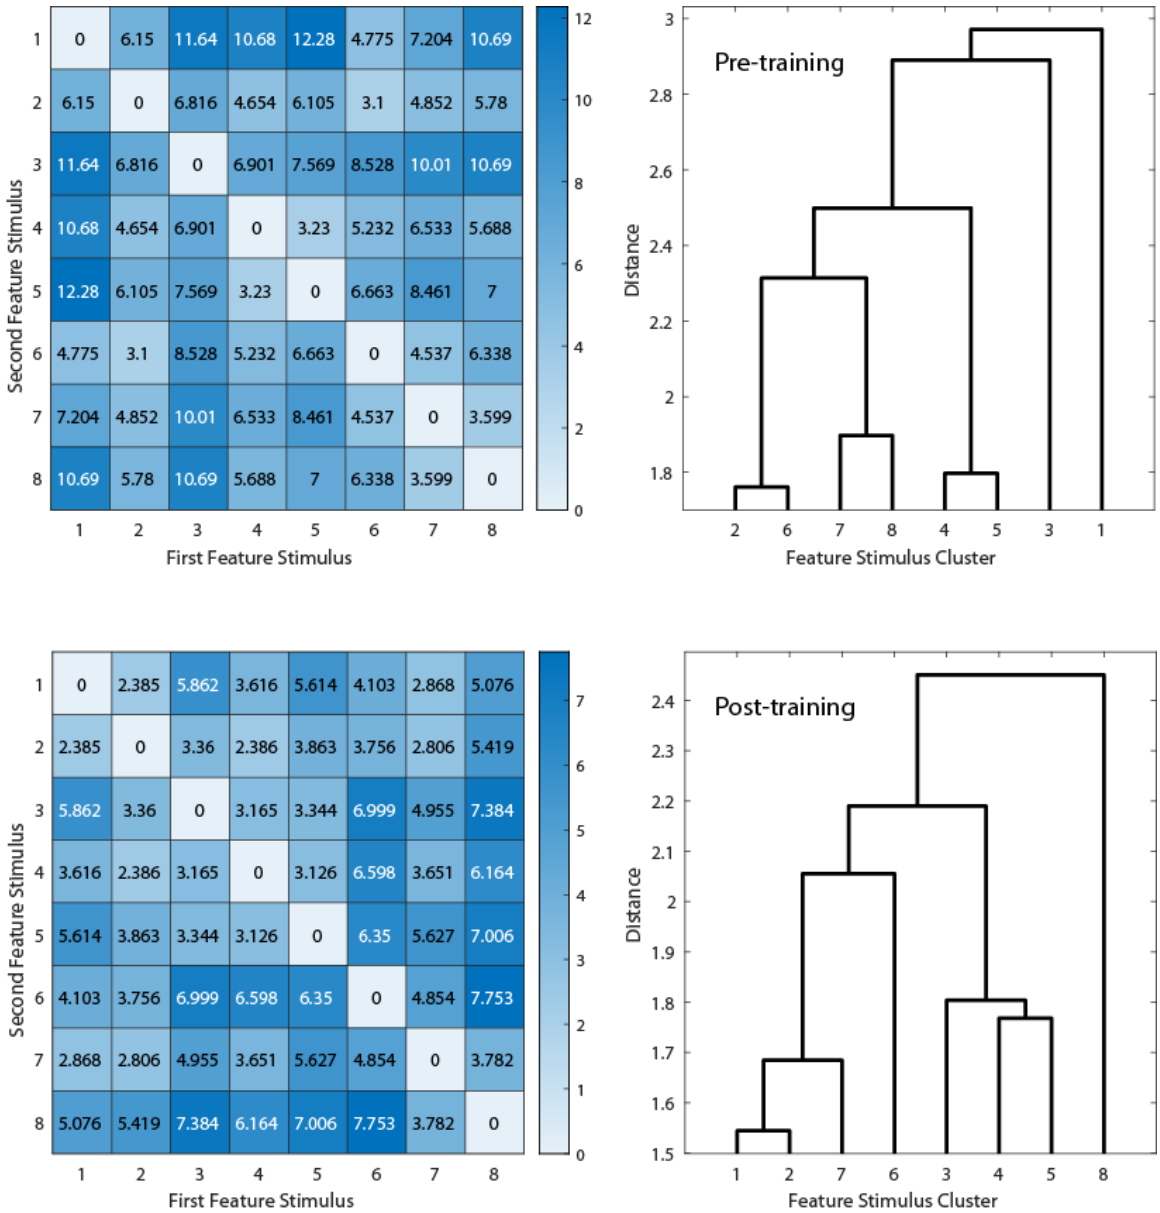

Figure 7. Heat map and dendrogram representing the representational distance between any two stimuli used in the feature task, before (top) and after training (bottom). Conventions are the same as in Fig. 4.