Task-specific topographic maps of neural activity in the primate lateral prefrontal cortex

Jinkang Derrick Xiang*[1,2], Megan Roussy[2,3], Benjamin Corrigan[2,4], Roberto A. Gulli[2,5], Rogelio Luna[2,6], Maryam Hasanzadeh Mofrad[1,7], Lyle Muller[1,7], Jörn Diedrichsen[1,8,9], Taylor W. Schmitz[1,2,3], Julio Martinez-Trujillo[2,3], Marieke Mur*[1,8,10]

[1]Western Institute for Neuroscience, University of Western Ontario, London, ON, Canada
[2]Robarts Research Institute, University of Western Ontario, London, ON, Canada
[3]Department of Physiology and Pharmacology, University of Western Ontario, London, ON, Canada
[4]Department of Biology, York University, North York, ON, Canada
[5]Zuckerman Mind Brain Behavior Institute, Columbia University, New York, The United States
[6]Facultad de Medicina y Ciencias Biomédicas, Universidad Autónoma de Chihuahua, Chihuahua City, Mexico
[7]Department of Mathematics, University of Western Ontario, London, ON, Canada
[8]Department of Computer Science, University of Western Ontario, London, ON, Canada
[9]Department of Statistical and Actuarial Sciences, University of Western Ontario, London, ON, Canada
[10]Department of Psychology, University of Western Ontario, London, ON, Canada

Abstract

Neurons in the primate lateral prefrontal cortex (LPFC) flexibly adapt their activity to support a wide range of cognitive tasks. Whether and how the topography of LPFC neural activity changes as a function of task is unclear. In the present study, we address this issue by characterizing the functional topography of LPFC neural activity in awake behaving macaques performing three distinct cognitive tasks. We recorded from chronically implanted multi-electrode arrays and show that the topography of LPFC activity is stable within a task, but adaptive across tasks. The topography of neural activity exhibits a spatial scale compatible with prior anatomical tracing work on afferent LPFC inputs. Our findings show that LPFC maps of neural population activity are stable for a specific task, providing robust neural codes that support task specialization. Moreover, the variability in functional topographies across tasks indicates activity landscapes can adapt, providing flexibility to LPFC neural codes.

Keywords: Lateral prefrontal cortex, rhesus macaque, multi-electrode electrophysiology, cognitive flexibility, functional maps, spatial scale

*Correspondence: jxiang27@uwo.ca (Jinkang Derrick Xiang), mmur@uwo.ca (Marieke Mur).

1. Introduction

Flexibility is one of the defining properties of higher-order cognitive functions supported by the primate lateral prefrontal cortex (LPFC). Unlike neural populations in primary sensory areas, where activity is dominated by stimulus features such as frequency and orientation [1, 2], LPFC neurons flexibly adapt their activity according to rules, behavioral context and feedback associated with different tasks, even when stimulus inputs are held constant [3, 4, 5, 6, 7]. This flexibility is shaped in part by the diversity of connections LPFC neurons receive, which is more heterogeneous than for sensory neurons [8, 9]. Because of their diverse connections, response profiles of LPFC neurons often exhibit selectivity to mixtures of task features, e.g., firing maximally only to a specific combination of rule, context and feedback [10, 11].

The selectivity of individual LPFC neurons to different combinations of task features creates unique challenges to understanding their functional organization as compared to sensory regions. In the visual cortex, for instance, individual neurons with preferences for similar stimulus features typically assemble into locally connected populations, giving rise to functional organizations which can be spatially delineated using stimulus mapping techniques [12]. Stimulating different visual field locations [13], stimulus orientations [14], or object categories [15, 16] reveals maps of neural populations with distinct feature preferences in striate and extrastriate cortex, respectively. The response profiles of these stimulus-tuned populations persist over days, weeks and months [17, 18, 19], forming stable topographic maps. The maps are organized at a columnar spatial scale [1, 20] that is detectable using recording techniques such as multi-electrode arrays and functional magnetic resonance imaging (fMRI) [13, 21, 22]. However, it remains an open question whether the LPFC exhibits task-specific functional topographies [23, 24, 25, 26, 27], and if so, whether these topographies are stable over time [28, 29], and have a spatial scale similar to those of other cortical areas. Testing for task-specific topographies of LPFC activity in primates, i.e., 'task mapping', has proven challenging due to constraints on the complexity and diversity of the 'task space' typically sampled in a given experiment [30]. Primate electrophysiology studies often probe only a single task, or task features which are not sufficiently distinct from one another, and do not analyze session-to-session variability in neural activity to characterize population stability.

If stable task-specific topographies exist in the LPFC at a spatial scale similar to sensory cortical areas, then this functional organization should be reliably detectable when mapping responses to different task features embedded in a sufficiently diverse task space. Here, we examined this possibility by acquiring a unique dataset consisting of multi-electrode array recordings from the LPFC of two macaques who were trained to perform three distinct cognitive

tasks [31, 32, 33]. The tasks cover a diverse task space by differentially recruiting a series of cognitive demands, embedding stimuli in rich behavioral contexts, and leveraging virtual reality environments for strong engagement of the LPFC [34]. The monkeys performed the same set of tasks over test sessions spanning multiple days, allowing us to assess the stability of task-tuned population responses over time.

To test for task-specific topographies in the LPFC, we take an approach inspired by representational similarity analysis (RSA) [35]. In fMRI [36, 37] and electrophysiology studies [38, 39], RSA is typically used to examine the distributed spatial pattern of responses across multiple measurement channels (e.g., voxels, electrodes) to an individual stimulus. Comparing these patterns across different stimuli allows one to infer representational geometry in the cortex, abstracted from the measurement channels themselves [40] (Fig. 1). Here, we adapt RSA to topographic similarity analysis (TSA) by focusing on the matrix transpose, i.e., the response profile of individual measurement channels (electrodes) to multiple task features [41]. Comparing these patterns across different channels allows us to infer their functional topography in the cortex, abstracted from the specific task features (Fig. 1).

TSA has two key advantages over stimulus-feature mapping strategies typically used for visual cortex [12]. First, TSA flexibly accommodates the full tuning profile of each channel to multiple task features, enabling the modeling of potentially context-dependent mixtures of neural selectivities. Second, because TSA abstracts from task features, it provides a characterization of LPFC functional topographies that can be directly compared across different tasks.

We show that topographies of LPFC activity are task-specific and stable within a task. We then demonstrate a spatial scale of functional organization consistent across all three task-specific functional topographies, which recapitulates prior anatomical tracing work examining the afferent input patterns of the LPFC from ipsilateral associational cortices and contralateral LPFC [42, 43, 44]. Our results indicate that the functional organization of the LPFC exhibits stable topographies of task-specific population activity, likely reflecting distinguishable mixtures of afferent sensory and cognitive input.

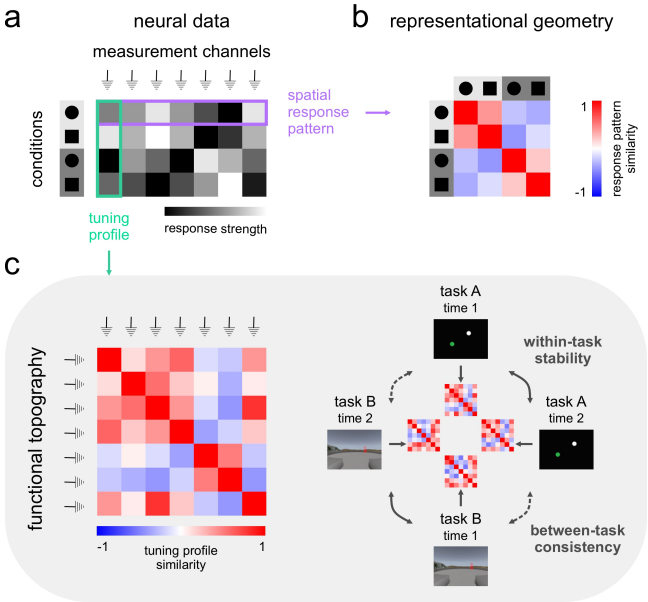

Figure 1: Topographic similarity analysis.

To test for task-specific functional topographies in the LPFC, we adapt representational similarity analysis (RSA) to topographic similarity analysis (TSA).

a) Simulated neural data across experimental conditions (rows) and measurement channels (columns). Higher spike rates in a given channel are denoted by lighter shades.

b) RSA abstracts from measurement channels to infer representational geometry.

c) TSA abstracts from experimental conditions to infer functional topography. TSA allows for quantitative comparison of functional topographies across time and tasks.

2. Results

Estimating the topography of LPFC neural activity across time and tasks

The same two rhesus macaque monkeys (monkey B and monkey T) each performed three cognitive tasks (Fig. 2a-c). The first task was an oculomotor delayed response task (ODR) [31]. The second task was a visuospatial working memory task (VWM) deployed in a virtual-reality environment with naturalistic scenes for stronger attentional engagement [32, 34]. The third task was a context-dependent decision making task (CDM), which was also deployed in a virtual-reality environment [33]. Altogether, these three tasks engage a wide spectrum of cognitive functions, including working memory, visuospatial attention, context-dependent decision making and motor planning.

We recorded the responses of neurons in layers II/III of LPFC areas 8A and 9/46, dorsal and ventral to the principal sulcus, using 96-channel multi-electrode Utah arrays (see Supplementary Fig. 1 for array placement). Each array covered a 4 mm × 4 mm cortical area with 10 × 10 electrodes (∼0.4 mm spacing). Both monkeys performed multiple measurement sessions for each task on separate days (Supplementary Fig. 2). In each session, we simultaneously recorded neural activity from multiple channels on each array (see Supplementary Table 1 for

details on active channels). Action potential times were extracted and synchronized to task events. To enable spatial mapping of response preferences across the array, we pooled spiking activity from putative units identified on the same recording channel, summing their activities to create multi-unit activity. Units measured by the same channel exhibited similar response preferences (Supplementary Fig. 3). The pooled activity reflects the activity of subpopulations of neurons within the area covered by each array channel.

To characterize the spatial organization of population codes in the LPFC, we first computed task tuning profiles, one for each channel in each session in each task (Fig. 2d-f). Task tuning profiles are vectors that store the firing rates of a channel to the experimental conditions. The time windows for estimating spike rates vary from a few hundred milliseconds up to a few seconds (see Methods for details). These time windows were determined based on task structure, monkey behavior and population decoding results. Next, we computed a channel-by-channel tuning similarity matrix for each array in each session in each task (Fig. 2g). Each element of the matrix represents the Pearson correlation of tuning profiles between a channel pair. The matrix as a whole reflects the similarity of task tuning for all channel pairs, thus capturing the functional topography of task-tuned LPFC activity. This characterization enables quantitative comparison of functional topographies across time and tasks (Figs. 1, 2g). For comparative purposes, we also analyzed trial-to-trial fluctuations around the trial averages that define the tuning profiles (see Methods). Topographies based on these spontaneous fluctuations, or residuals, are expected to be consistent across tasks [45, 41].

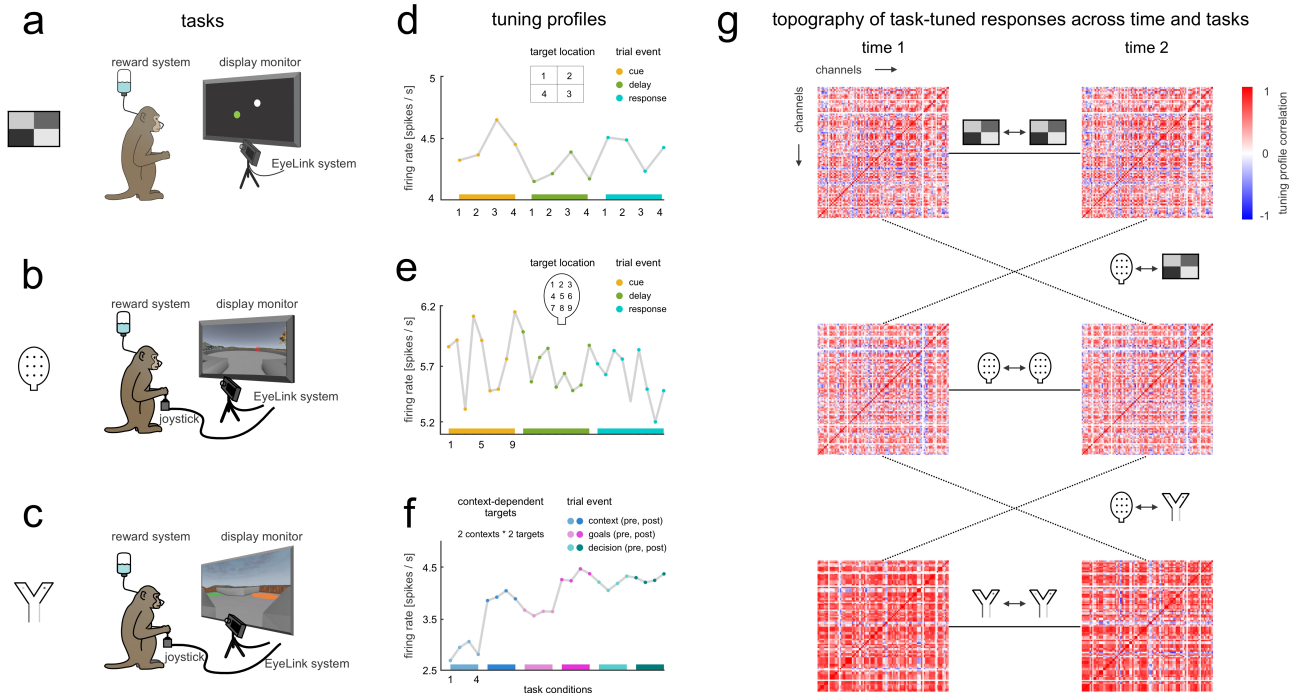

Figure 2: Comparing task-tuned functional topographies in the LPFC across time and tasks. a) Oculomotor delayed response task (ODR). Monkeys fixated a point on the screen. A visual cue appeared, then disappeared. After a delay, the monkeys saccaded toward the remembered target location.

b) Visuospatial working memory task (VWM). A visual cue appeared in one of nine target locations in a virtual arena, then disappeared. After a delay, the monkeys navigated toward the remembered target location using a joystick.

c) Context-dependent decision making task (CDM). Monkeys navigated through an X maze using a joystick. The texture of the corridor walls indicated the decision context for target selection.

d-f) Example tuning profile for a single channel in the ODR (d), VWM (e), and CDM (f) tasks. Legends indicate the task features that composed the unique task space.

g) Schematics for TSA of task-tuned responses across time (solid lines) and tasks (dashed lines). Matrix elements are Pearson correlations of tuning profiles between channel pairs.

The topography of LPFC neural activity is stable over time but adaptive across tasks

LPFC neural activity was on average positively correlated between channels, with stronger correlations observed for task-tuned than residual responses. In the tuning similarity matrices, the mean cross-channel correlation (r) over sessions and tasks was 0.50+/-0.21 (mean+/-SD) for monkey B (mB) and 0.44+/-0.16 for monkey T (mT). In the residual similarity matrices, the mean cross-channel correlation (r) over sessions and tasks was 0.08+/-0.04 for mB and 0.15+/-0.08 for mT. The sign and magnitudes of the observed correlations are consistent with prior work [5, 46, 47].

We next assessed whether the functional topography of LPFC neural activity is consistent over time (sessions) and across tasks (ODR, VWM and CDM). To do so, we used TSA to

compute the average correlation of the channel-by-channel similarity matrices across sessions within and between tasks (Fig. 2g). Session-to-session correlations within the same task assess the consistency of topography over time, and session-to-session correlations between different tasks assess the consistency of topography across tasks. Consistencies were assessed for topographies based on task-tuned responses and residuals separately, for each array in each monkey. To control for array shifting across time, we only included sessions spaced apart no more than 20 days for within- as well as between-task comparisons. Given that data for some tasks were acquired more than 20 days apart (see Supplementary Fig. 2), between-task comparisons are based on two out of three tasks for each monkey. We focus on results from the ventral array for both monkeys in the following sections, due to low signal-to-noise ratio of dorsal array recordings in mT (see Methods).

Over time, the task-tuned topographies are relatively stable. The consistencies of task-tuned topographies were as follows: mB: VWM-VWM $r = 0.43$, mB: CDM-CDM $r = 0.25$, mT: VWM-VWM $r = 0.64$ and mT: ODR-ODR $r = 0.21$. In all cases, the r values for within-task consistencies were significantly larger than zero (one-sided permutation test; $p < 0.05$, Fig. 3a,b). Moreover, the within-task consistencies for task-tuned topographies were significantly higher than those observed for residual topographies (two-sided t-test; mB: VWM-VWM $t(54) = 3.97$, $p < 0.001$, mB: CDM-CDM $t(14) = 2.95$, $p = 0.01$, mT: VWM-VWM $t(7) = 5.79$, $p < 0.001$ and mT: ODR-ODR $t(5) = 2.60$, $p = 0.05$). The dorsal array in mB showed similar results (Supplementary Fig. 4). Hence, over sessions of a given task, the topographies of task-tuned responses are more consistent than the topographies of concurrent trial-to-trial fluctuations in spontaneous activity.

Across different tasks, the task-tuned topographies adapt. The consistencies of task-tuned topographies were as follows: mB: VWM-CDM $r = 0.11$, mT: VWM-ODR $r = 0.18$. The r values for between-task consistencies were significantly larger than zero (one-sided permutation test; $p < 0.05$). However, they were significantly below their noise ceilings estimated from the within-task consistencies (one-sided t-test; mB: VWM-CDM $t(36) = -18.04$, $p < 0.001$, mT: VWM-ODR $t(3) = -7.25$, $p = 0.003$, Fig. 3c,d), suggesting that task-tuned functional topographies are not fully consistent across tasks even when considering noise inherent to the data. Moreover, the r values quantifying between-task consistency of task-tuned topographies were either significantly lower than or not significantly different from those observed for residual topographies (two-sided t-test; mB: VWM-CDM $t(36) = -3.03$, $p = 0.005$, mT: VWM-ODR $t(3) = -0.51$, $p = 0.64$). The dorsal array in mB showed similar results (Supplementary Fig. 4). Hence, between a given pair of two different tasks, the topographies of task-tuned responses

are less consistent than the topographies of concurrent spontaneous trial-to-trial fluctuations in activity.

The above analyses average correlations across session pairs to estimate the overall consistency of the LPFC topography across time and tasks. While this approach provides a general measure of stability, it does not capture finer, continuous changes in topography over time or quantify the relative contributions of time and task to topographic change. To address this, we conducted an additional analysis using linear regression to quantify the unique contributions of time and task to explaining variance in session-pair consistency (Supplementary Results 1, Supplementary Fig. 5a). Across sessions and arrays, a linear regression model including predictors for both time and task explained on average 87% of the variance in session-pair consistency (Supplementary Fig. 5b). Time uniquely accounted for 15% of the explainable variance, while task explained 40%, further underscoring the adaptability of the LPFC topography across tasks relative to its stability over time (Supplementary Fig. 5c). Although variability is present across arrays, these results support the distinct roles of time and task in shaping LPFC topography.

To evaluate the robustness of our approach in identifying consistent and meaningful topographic patterns, we tested its sensitivity to different low-dimensional projections of the data using a split-half approach with non-overlapping conditions. This analysis tested whether the within-task consistency of tuning similarity matrices remains stable when derived from disjoint subsets of conditions. For both same-session (Supplementary Fig. 6a) and cross-session (Supplementary Fig. 6b) comparisons, within-task consistencies remained significantly larger than zero (one-sided t-tests, $p < 0.05$). These findings support the reliability of our approach in capturing stable and meaningful topographic patterns despite variations in condition sampling. Importantly, the between-task consistencies observed in this control analysis reproduced the original findings (Fig. 3c,d): they were significantly below the noise ceiling estimated from the within-task consistencies (Supplementary Fig. 6b; one-sided t-tests; mB ventral: VWM-CDM $t(36) = -15.04$, $p < 0.001$, mB dorsal: VWM-CDM $t(36) = -43.68$, $p < 0.001$, mT ventral: VWM-ODR $t(3) = -5.08$, $p = 0.007$), reinforcing the adaptability of the LPFC functional topography across tasks.

Our analyses suggest that task-tuned functional topographies in the LPFC are (1) stable across time within a task: neural populations with similar tuning on one day tend to exhibit similar tuning on another day, and (2) adaptive across tasks: neural populations with similar tuning in one task do not necessarily exhibit similar tuning in another task. These findings demonstrate the capacity of LPFC to maintain consistent neural population topographies over

time while flexibly adapting to the demands of different tasks.

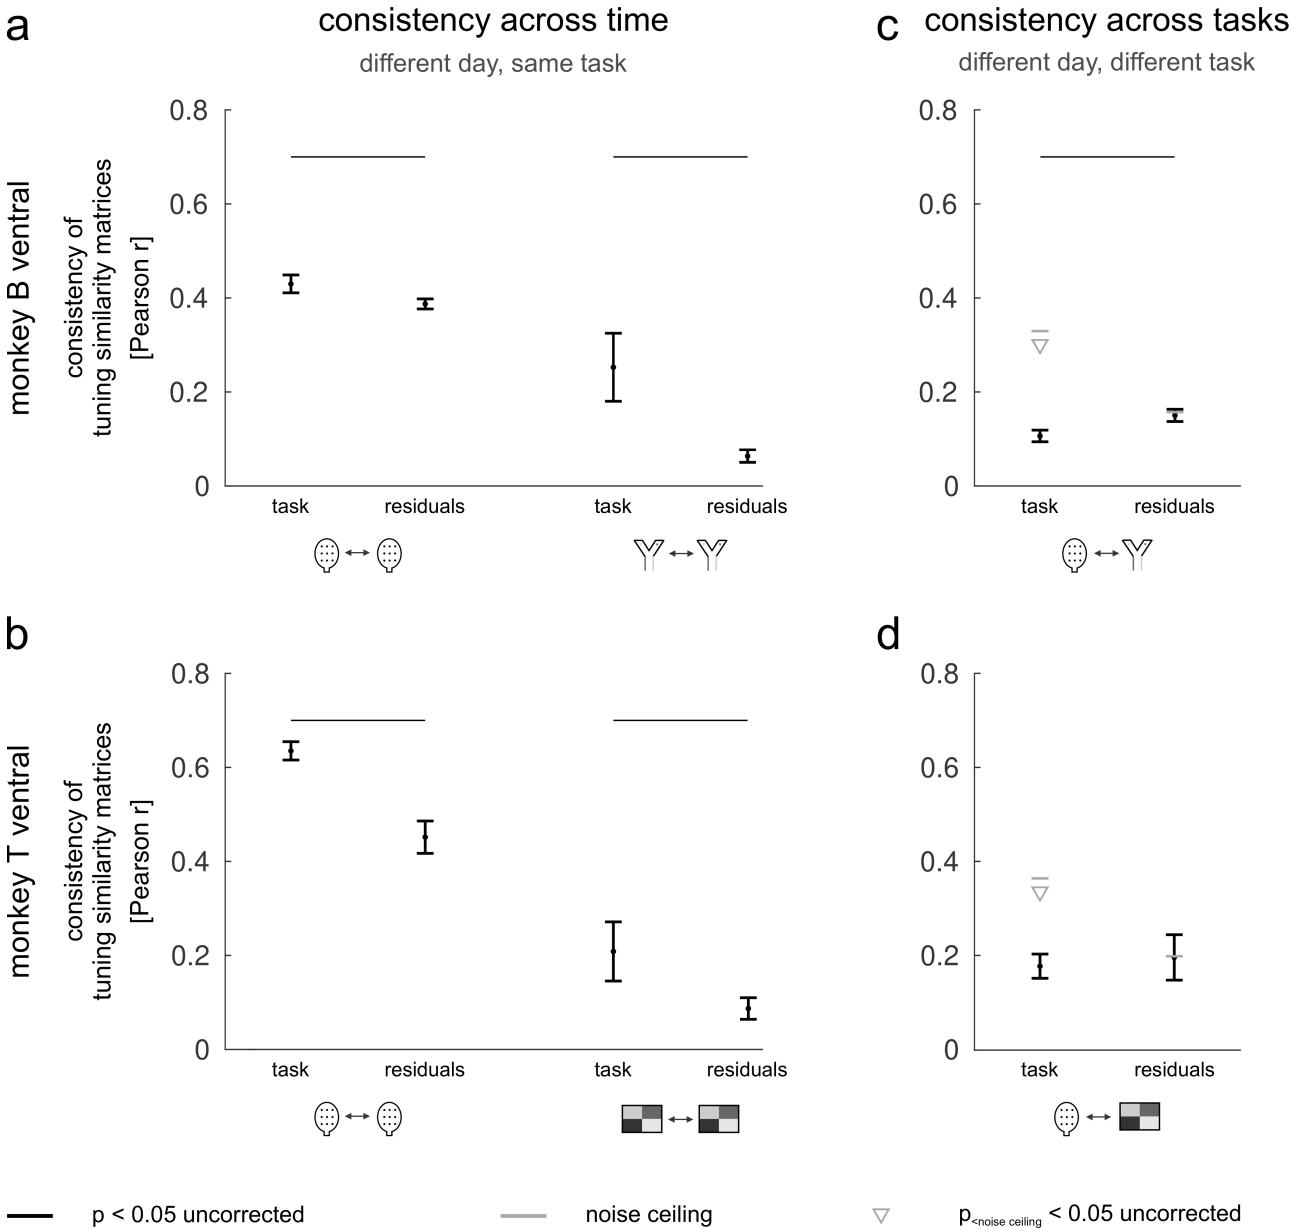

Figure 3: Task-tuned functional topographies in the LPFC are stable across time but adaptive across tasks.

a-b) Consistency of functional topographies across time within a task for mB (a) and mT (b).

c-d) Consistency of functional topographies between tasks for mB (c) and mT (d).

Black horizontal lines indicate significant differences between task and residuals ($p < 0.05$, two-sided paired t-test). Gray horizontal lines show the noise ceiling. Gray triangles indicate values significantly below the noise ceiling ($p < 0.05$, one-sided t-test). Error bars show the standard error of the mean (SEM) across session pairs.

Linking the topography of LPFC neural activity to fine-grained spatial maps

Given that feature-tuned neurons are known to cluster in populations at the spatial scale of cortical columns [1, 2, 20, 25], the observed task-specific functional topographies in the LPFC (Fig. 3; Supplementary Fig. 4) may be expressed spatially by clusters of similarly tuned neurons with an adaptive task-dependent 'fingerprint'. To quantify the degree of spatial clustering

of channels with similar response preferences on the array, we computed spatial autocorrelation functions (ACFs) of channel tuning profiles [27, 26, 48, 49]. The ACFs measure the tuning similarity between channels as a function of their distance on the array. We computed ACFs for each session and task. As a measure of spatial autocorrelation, we used a cross-validated implementation of global Moran's I that is bounded between -1 and 1, where positive values indicate similar tuning (see Methods). We found that the spatial autocorrelation between channels decreases exponentially as their distance on the array increases (Fig. 4a; Supplementary Fig. 7a,b). From the ACFs of both monkeys, we observed a positive spatial autocorrelation up to a distance of 1-2 mm across tasks (Fig. 4a).

These observations suggest that the task-specific functional topographies in the LPFC converge on a common spatial scale. To quantify this spatial scale, we fit Laplacian functions to the observed ACFs. Modeling the ACFs with Laplacians provides a concise description of the observed spatial decay using a limited set of parameters, allowing us to derive a scalar estimate of spatial scale. This estimate, represented by the full-width-at-half-maximum (FWHM) of the fitted Laplacians, captures the spatial extent of similarity in tuning. The FWHM is interpreted as the width of the 2-dimensional kernel needed to smooth spatially independent tuning profiles to yield the same degree of clustering we observe in the data [50]. The median of the FWHMs across sessions and tasks is 369 +/- 107 μm (SD), suggesting the existence of fine-grained, or mesoscale, clusters of similarly tuned neural populations on the arrays.

To get an impression of the spatial structure of the tuning similarity on the arrays, we next visualized the data using 2-dimensional (2D) multidimensional scaling (MDS). We converted the tuning profile correlations to correlation distances, and applied 2D MDS. Channels were color-coded based on their location in the 2D MDS space and projected back to the arrays [41] (Fig. 4b). Each color indicates a unique channel tuning profile across all conditions of a given task. Channels with similar tuning profiles are thus similarly colored, and their relative position on the array reveals observable clustering, which is highly similar within tasks and moderately to weakly similar between tasks. Across all maps, the clusters of similarly task-tuned channels reveal a consistent spatial scale (Fig. 4c; Supplementary Fig. 7c,d). To further emphasize the spatial scale of these clusters, we smoothed the 2D MDS array maps using 2D Gaussian kernels with FWHM values matching the empirically derived cluster sizes from the fitted Laplacians. In all three tasks, this visualization highlights the mesoscale organization of task-tuned population activity, showing clusters at a spatial scale similar to that observed for maps in sensory cortical areas, such as V1, where mesocale functional organization reflects tuning to stimulus features [14, 21] (Fig. 4d).

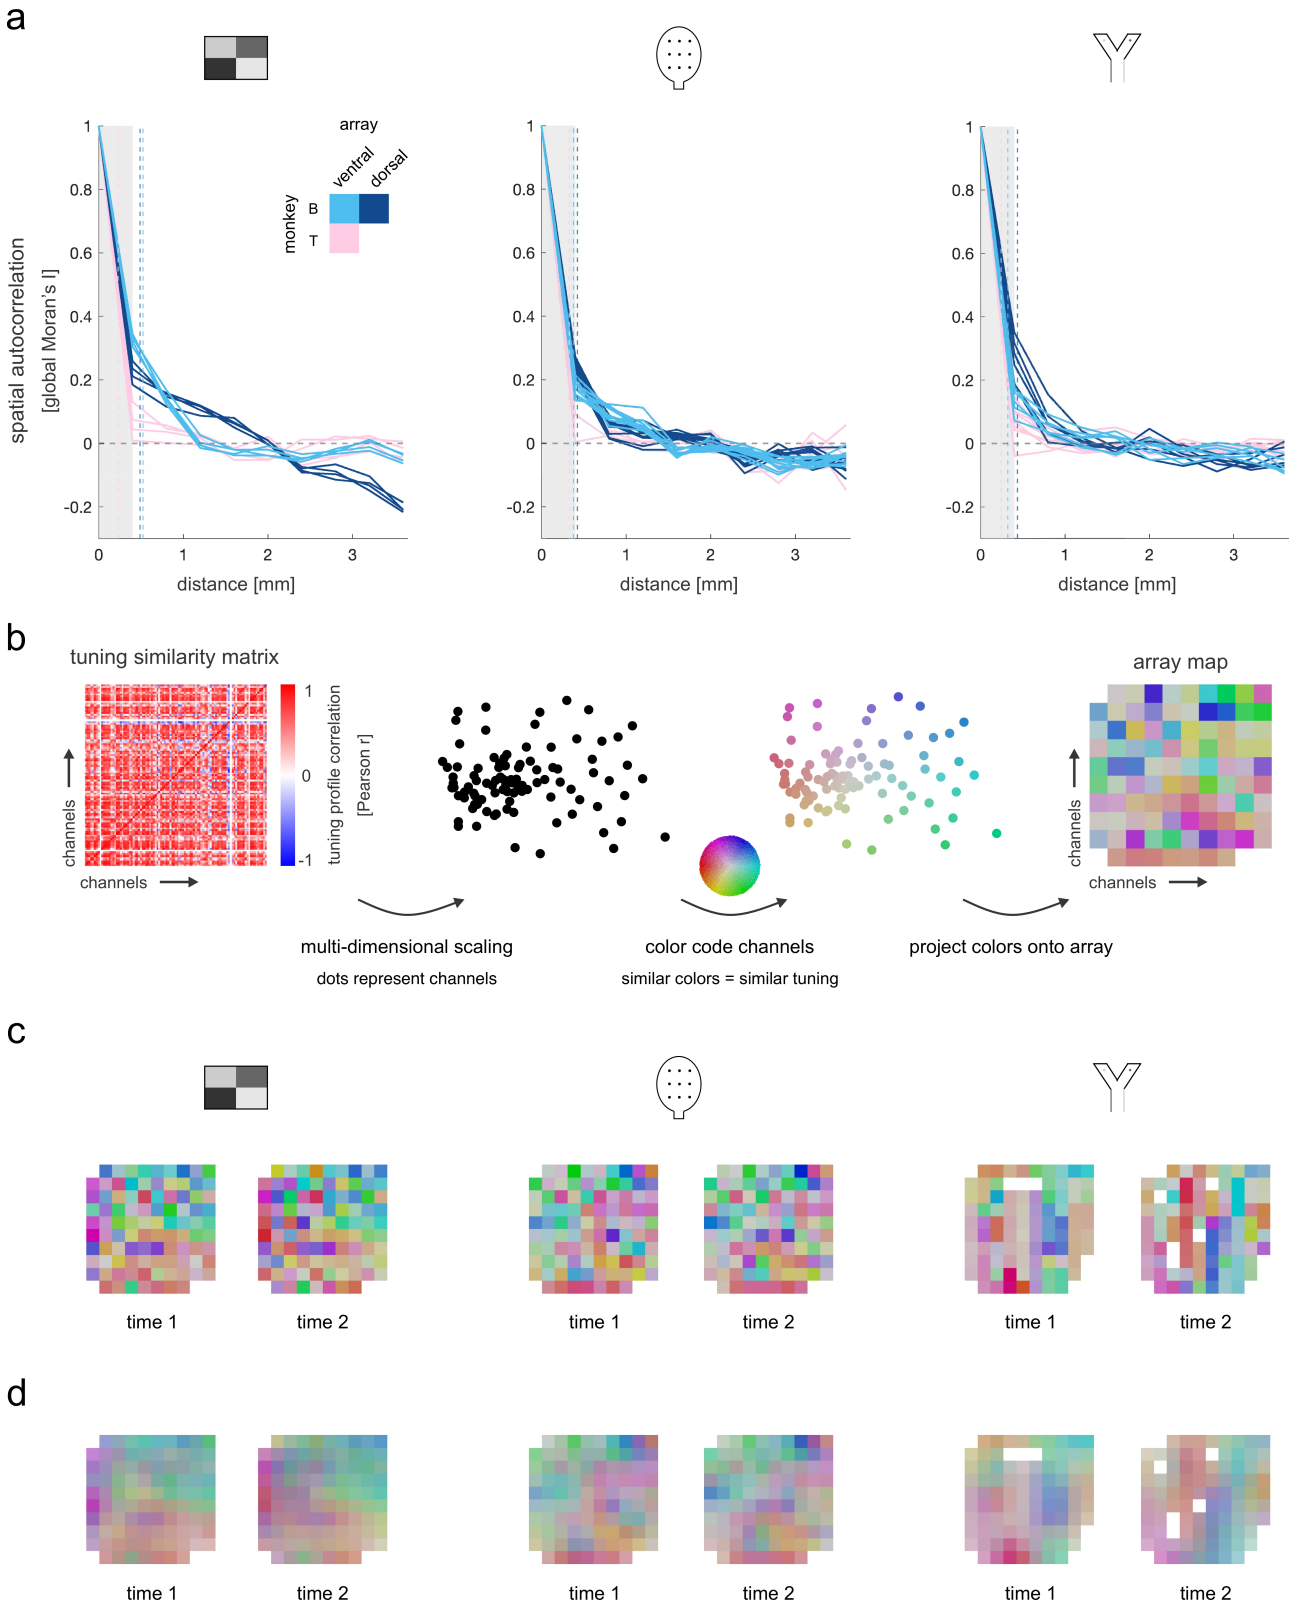

Figure 4: Task-tuned functional topographies in the LPFC are organized at a fine-grained spatial scale.

a) Spatial autocorrelation functions (ACFs) for ODR, VWM and CDM tasks. Solid lines represent the ACFs for individual sessions in each task. Colors indicate arrays: light blue for the mB ventral array, dark blue for the mB dorsal array, and pink for the mT ventral array. Gray shaded areas indicate the width between two immediately neighboring channels on the array (0.4 mm). Vertical dashed lines show the median full-width-at-half-maximum (FWHM) of Laplacian functions fit to ACFs of individual sessions.

Figure 4:

b) Steps taken to visualize tuning similarity on the array. Channels with similar tuning profiles in a given task are similarly colored.

c) Array maps for the mB dorsal array in all three tasks, using two example sessions per task.

d) Array maps in panel c after smoothing with 2D Gaussian kernels whose FWHMs match that of the fitted Laplacian functions.

Our findings suggest that task-tuned LPFC responses are spatially clustered at the mesoscale, similar to sensory cortical areas [1, 2, 20, 14, 21]. We next asked what properties of cortical organization might shape this spatial scale. One possibility is that the afferent input projections from different regions of the brain may exhibit a spatial scale of clustering in the LPFC similar to that of the functional topographies observed in the present work. This would suggest that the structural organization of long-range inputs into the LPFC shapes the functional topographies of its task-tuned responses. To test this hypothesis, we analyzed a well-documented structural connectivity map representing the profile of afferent inputs to the macaque principal sulcus of the LPFC from the contralateral principal sulcus through callosal fibres, and from the ipsilateral parietal cortex through associational fibres (Fig. 5a,b) [42, 51]. The structural map overlaps the cortical patch measured with multielectrode arrays in the current study (Fig. 5a,b). To obtain an unbiased estimate of the map's intrinsic spatial scale, we sampled 210 patches at random locations from the map, using a 4 mm $\times$ 4 mm patch size equivalent to the area covered by our arrays (Fig. 5b,c). For each sample, we computed a spatial ACF and corresponding FWHM using the same strategy as depicted in Fig. 4a. For both measures, we observed that the functionally derived estimates in Fig. 4a fall within the 95% confidence interval of the estimates computed across the structural map samples (Fig. 5d). These findings suggest that the structural and functional spatial scales are comparable, consistent with the idea that long-range afferent inputs may contribute to the mesoscale organization in the LPFC.

To assess the sensitivity of these results to stripe size in the structural map, we conducted additional simulations by scaling the map to produce smaller, original, and larger stripe sizes (Supplementary Fig. 8a-c). Larger stripe sizes produced broader spatial scales in the structural ACFs, positioning the functionally derived ACFs outside the 95% confidence interval of the structural distribution at shorter distances (Supplementary Fig. 8f). Results for the smaller stripe size were less conclusive (Supplementary Fig. 8d) but did not challenge the alignment observed for the original stripe size (Supplementary Fig. 8e). These findings reinforce the link between structural and functional spatial scales by demonstrating that a broader structural scale of long-range afferent inputs is inconsistent with the functional measures.

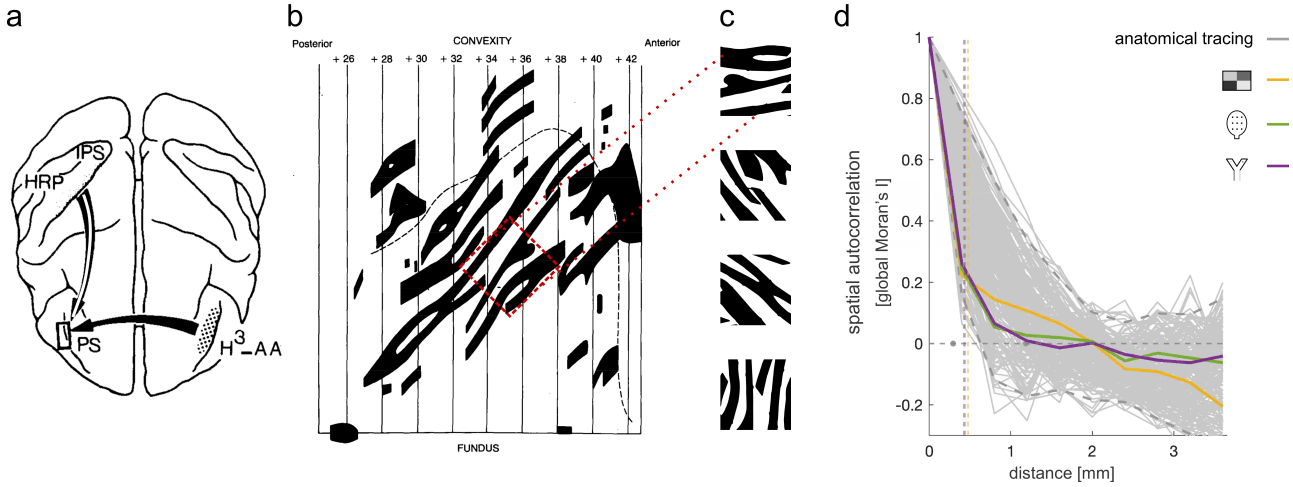

Figure 5: Linking task-specific functional topographies in the LPFC to fine-grained structural maps.

a) Anatomical tracing of white matter fiber inputs to the macaque LPFC [43, 42], adapted with permission from [42]. The LPFC receives inputs through contralateral callosal fibers of the principal sulcus (PS; solid black arrow) and ipsilateral association fibers from the intraparietal sulcus (IPS; white strike arrow). HRP = horeradish peroxidase, $H^3 - AA$ = tridared amino acids.

b) Structural map representing the spatial profile of afferent inputs to the macaque LPFC, adapted with permission from [42]. The black dashed line shows the rim of the PS. Black stripes show the reconstructed terminal fields of callosal fibers in the PS, reflecting inputs from the contralateral PS. Interdigitated white stripes receive inputs from the ipsilateral IPS [43]. The red square shows an example cortical patch sampled for our analysis.

c) Examples of cortical patches sampled at random locations from the structural map in panel b. Each sample has a coverage of 4 mm $\times$ 4 mm.

d) Functional ACFs superimposed on structural ACFs. Yellow, green and purple lines represent mean functional ACFs across sessions for the ODR, VWM and CDM tasks, respectively, for the mB dorsal array (Fig. 4a). Vertical dashed lines show the mean FWHM across sessions for each task. Solid gray lines represent structural ACFs for 210 structural map samples. Dashed gray lines show the 95% confidence interval of the structural ACFs. Gray dots placed on the zero autocorrelation line show the 95% confidence interval of the structural FWHMs.

3. Discussion

We used TSA to characterize the functional topography, temporal stability, and spatial scale of task-tuned LPFC neural activity by analyzing array recordings from awake behaving macaques performing three distinct cognitive tasks [31, 32, 33]. TSA revealed that the spatial topography of task-tuned LPFC neural activity is stable across time within a task but adaptive across tasks. The stability for task-tuned responses is higher than for concurrent spontaneous fluctuations, indicating that the correlation structure among LPFC neural populations is strengthened by task. We further demonstrate that although all three tasks exhibit distinct topographies of task-tuned activity, they converge on a common spatial scale. Finally, we show

that this spatial scale is likely shaped, in part, by the organization of afferent long-range inputs into the LPFC from distinct cortical areas throughout the brain.

Our findings indicate that TSA provides a powerful and flexible strategy for decomposing high-dimensional response profiles from LPFC array recordings into topographies of task-tuned preferences. However, TSA has some notable differences from traditional methods for deriving neural tuning profiles [12]. TSA infers topographic patterns indirectly, based on second-order correlations between channels rather than direct measures of tuning to isolated, explicitly defined task features [41] (Fig. 1). In this sense, the "maps" revealed by TSA are not traditional tuning maps anchored to a single feature dimension but rather spatial patterns of similarity in neural response profiles. This has multiple advantages: TSA enables the modeling of potentially context-dependent mixtures of neural selectivities and provides a characterization of functional topographies that can be directly compared across different tasks. Moreover, we show that the use of a noise ceiling is a helpful theoretical construct for contextualizing effect sizes in TSA [52]. However, by itself, TSA does not eliminate the fundamental challenges associated with disentangling the complex interplay of stimulus features, task structure, and intrinsic connectivity in shaping LPFC functional organization. Future approaches could integrate TSA with more refined task designs and multimodal imaging measures to yield an even clearer understanding of the spatial and functional organization of the LPFC in both macaques and humans. Below we discuss some of these methodological challenges in light of the current study.

Task mapping in the LPFC presents significant challenges. Unlike stimulus mapping experiments in the visual cortex, which are relatively easy to parameterize according to objective physical features such as location, orientation, or object category, task mapping experiments are more difficult to parameterize because task features often represent abstract concepts, such as rule associations or temporal interdependencies [3, 5, 10, 41, 26, 27]. For example, although the LPFC exhibits task-specific topographies, we also observed that these topographies still share a significant portion of task-related variance. We interpret this shared variance as tuning to shared features between tasks. For instance, all three tasks consist of visual stimuli and a trial sequence with successive cue and target phases. Future work examining task mapping in the LPFC will benefit from experimental designs which increase the number and diversity of task features sampled, and the precision of their relationship to one another in a high dimensional "task space" [30]. Task mapping experiments which systematically vary task-stimulus combinations in a high dimensional task space will be better equipped to explore the similarity of distinct tasks to one another, as well as their potential linear and nonlinear integration,

thereby enabling more computationally principled hypotheses about LPFC function, including biased competition [53, 54] and compositional coding [55, 56].

Connectomic profiling of the LPFC also presents significant challenges. In contrast to the relatively homogeneous short-range structural connections that shape topographic maps in the visual cortex [21], LPFC functional topographies are likely shaped by a heterogeneous mixture of long-range afferent inputs and local intrinsic connectivity [8, 57, 58, 59, 23, 60]. Anatomical tracing studies have shown that the LPFC receives patterned long-range input projections from multiple areas — most notably from the posterior parietal cortex and the contralateral LPFC — that form elongated stripes of connectivity [8, 58]. We estimated the spatial scale of LPFC afferent input patterns by sampling a structural connectivity map provided in seminal work by Patricia Goldman-Rakic [42]. We show that the spatial scale of this map corresponds closely to the ∼300–400 μm scale of the functional topographies observed here, despite the vastly different methodologies used to derive the structural and functional maps. The stripes, along with local horizontal connections that form patchy networks [59] may serve as substrates for the spatially clustered functional responses to task features. Recent multimodal studies of the macaque LPFC have begun to bridge the structure-function gap using microstimulation [57] and diffusion MRI [61]. Consistent with our findings, these studies indicate that a diverse network of long-range structural projections may provide the scaffolding for LPFC's functional organization, influencing the intrinsic network structure upon which task-specific activity patterns unfold [62, 57]

The spatial scale of the LPFC functional topographies observed in macaques suggests that these patterns could, in principle, be resolved using TSA in combination with ultra-high-field fMRI in humans. While task-based fMRI studies of human LPFC have not yet achieved such fine-grained resolution, there is precedent in the visual domain, where stimulus-driven ocular dominance columns, organized at submillimeter scales, have been successfully visualized with cutting-edge imaging techniques (e.g., [21]). Future cross-species translational efforts, leveraging the complementary strengths of human fMRI and macaque electrophysiology with a unified analytical framework such as TSA, may achieve a more comprehensive understanding of LPFC function. In humans, the non-invasive nature of fMRI allows exploration of a wider and more diverse task space, encompassing complex and abstract cognitive functions that are challenging to systematically manipulate in macaque models. The rich datasets obtained from human fMRI studies may generate novel hypotheses regarding the functional organization and computational principles of the LPFC, which can be reverse translated through targeted electrophysiological experiments in macaques to provide mechanistic insights into the observed

mesoscopic patterns.

In summary, our application of TSA to LPFC recordings in awake macaques has elucidated stable yet task-specific spatial topographies of neural activity, all converging on a consistent spatial scale likely influenced by long-range cortical inputs. These results demonstrate TSA's efficacy in decomposing complex neural response patterns into spatially interpretable topographies and highlight the nuanced interplay between task demands and structural connectivity in shaping LPFC function. Furthermore, the alignment of macaque functional topographies with the potential resolution of ultra-high-field fMRI in humans paves the way for cross-species translational research. By integrating the expansive and diverse task spaces accessible through human fMRI with the precise electrophysiological insights from macaque models, this unified analytical framework promises to deepen our understanding of the LPFC's organizational principles and its role in executive cognitive processes.

4. Methods

4.1. Subjects and ethics statement

We recorded LPFC neural activity in two male rhesus macaques (*Macaca mulatta*, monkey B and monkey T, 10 and 9 years old) while they were performing three different cognitive tasks across multiple measurement sessions. All training, surgery, and recording procedures conformed to the Canadian Council on Animal Care guidelines and were approved by The University of Western Ontario Animal Care Committee.

4.2. Behavioral tasks

4.2.1. Oculomotor delayed response (ODR) task

Fig. 2a illustrates the experimental setup of the oculomotor delayed response task. Each trial began with the appearance of a fixation point at one of 16 predefined locations on a computer screen. Then a target was presented for 1000 ms before it disappeared. The monkeys were asked to maintain fixation for a variable length delay period (1400 ms - 2500 ms, median = 1800 ms) and upon extinction of the fixation point, make a saccade towards the remembered target location to get a reward. More information on the ODR task can be found in [31].

4.2.2. Visuospatial working memory (VWM) task

Fig. 2b illustrates the experimental setup of the visuospatial working memory task, which took place in a virtual reality environment. Within the virtual arena in the environment, targets were arranged in a 3×3 grid. The time needed to navigate between adjacent targets

was ∼0.5 s. During the cue period, a visual cue (red rectangle) was presented in one of the nine target locations for 3 seconds, then disappeared. After a 2-second delay, the monkeys navigated towards the remembered target location at a constant speed using a joystick. Upon reaching the correct target, the monkeys received a reward. More information on the VWM task can be found in [32].

4.2.3. Context-dependent decision making (CDM) task

Fig. 2c illustrates the experimental setup of the context-dependent decision making task, which was also deployed in a virtual reality environment. The task took place in a double ended Y maze, also termed "X maze" as in [34, 63]. The monkeys navigated through the X maze using a joystick. The texture of the walls, being brown "wood" or dark grey "steel", indicated which coloured disk the monkeys should choose at the bifurcation to get a reward. In other words, the decision context was specified by the texture of the walls. We will refer to the coloured disks as goals in subsequent text. More information on the CDM task can be found in [33].

4.3. Neural recordings

Two 96-channel Utah arrays (4 mm × 4 mm coverage, 10 × 10 electrodes, spaced at 0.4 mm, 1.5 mm in length) (Blackrock Microsystems) were chronically implanted in the left LPFC in each animal. They were located anterior to the arcuate sulcus and on the ventral and dorsal side of the posterior end of the principal sulcus, targeting layers II/III of cortical areas 8A and 9/46 (see Supplementary Fig. 1 for array placement).

Neural data were recorded using a Cerebus Neural Signal Processor (Blackrock Microsystems). The neural signal was digitized (16 bit) at a sampling rate of 30 kHz. Data were online sorted to identify putative units on each recording channel in real-time. Action potential times were extracted and synchronized to task events. For all tasks, eye positions were monitored using SR Research EyeLink 1000, at a sampling rate of 500 Hz.

Both monkeys performed multiple measurement sessions for each task. Sessions were acquired on separate days. Monkey B performed 4 sessions for the ODR task, 11 sessions for the VWM task, and 6 sessions for the CDM task. Monkey T performed 4 sessions for the ODR task, 8 sessions for the VWM task, and 9 sessions for the CDM task. Three sessions were excluded from analysis in the VWM task for monkey T due to low trial numbers across experimental conditions. See Supplementary Figure 2 for further details on the included sessions. The dorsal array in monkey T was excluded from analysis due to low signal-to-noise ratio of recordings at the time of the experimental sessions, which makes it difficult to map the topography of the array.

Quality assurance data generated by the Blackrock software before each recording session indicated potential crosstalk in 4 out of the 60 sessions analyzed across monkeys and arrays. This affected only one monkey, array, and task (monkey B, dorsal, CDM) in 4 of the 6 measurement sessions for that task. To assess the robustness of the spatial scale results (Fig. 4) to potential crosstalk, we performed a conservative control analysis. For the affected sessions, we removed any shared spikes for channel pairs with more than 10% shared spikes, which impacted approximately 50 channel pairs on average per session. Results from this analysis confirmed that spike removal had minimal effects on the estimated spatial scale (Supplementary Fig. 7) and did not affect the pattern of results for the consistency analyses reported in Supplementary Figure 4.

4.4. Neural data analysis

4.4.1. Computing task tuning profiles

To enable spatial mapping of response preferences across the array, we pooled units recorded from the same channel by summing up their spiking activities. Units recorded from the same channel exhibited similar response preferences (Supplementary Fig. 3). We refer to the pooled activity as multi-unit activity. The pooled activity reflects the activity of subpopulations of neurons within the area covered by an array channel. This area has an estimated diameter of 300 microns based on the impedance of the electrodes.

For each channel in each session in each task, we first computed trial-specific spike rates for each experimental condition and applied a square root transformation to the spike rates to account for the Poisson-like increase of variability with increasing mean firing rates [64, 48]. We next computed trial-averaged spike rates by averaging across trials of the same experimental condition, which yielded a task tuning profile for each channel.

More specifically, in the ODR task, the tuning profile for a channel was defined as the trial-averaged spike rates for the 4 quadrants (16 targets were grouped into 4 quadrants based on their location in the retinotopic reference frame, labelled 1-4 starting from bottom left, clockwise) during cue (1000 ms), delay (1400 ms - 2500 ms, median = 1800 ms), and response (first 500 ms) epochs. In the VWM task, the tuning profile for a channel was defined as the trial-averaged spike rates for the 9 target locations during cue (3000 ms), delay (2000 ms), and response (first 500 ms) epochs. In the CDM task, the tuning profile for a channel was defined as the trial-averaged spike rates for the combinations of decision contexts and goal configurations in time windows before and after context onset (50 ms before, 600 ms after), goals onset (500 ms before, 300 ms after), and decision onset (500 ms before, 500 ms after). Goal configuration refers to the location of the disk colour that is associated with wood. If the colour associated

with wood is on the left-hand side at the bifurcation, the trial is in configuration 1, otherwise configuration 2.

We then subtracted out the mean firing rate across all experimental conditions. The mean-centered tuning profile reflects the modulation in firing rate by the experimental conditions, providing a rich characterization of channel tuning in each task. The time windows used for computing tuning profiles were based on trial structure, monkey behavior, and neural population decoding results. Results of subsequent analyses do not critically depend on the exact time windows used.

4.4.2. Computing the topography of task-tuned activity

For each array in each session in each task, we computed a channel-by-channel tuning similarity matrix. Elements in this matrix reflect the Pearson correlation of tuning profiles between channel pairs. The matrix reflects the similarity of task tuning for all channel pairs, thus capturing the similarity structure of tuning across the array. As such, it provides the basis for mapping of tuning similarity across the cortical sheet. This approach can be used to study functional topography [41, 65]. Importantly, the tuning similarity matrices abstract from the specific experimental conditions used in a single task, which enables comparison of functional topographies between tasks.

4.4.3. Computing the topography of spontaneous activity

For comparative purposes, we also analyzed trial-to-trial fluctuations about the trial averages that define the tuning profiles. Topographies based on these spontaneous fluctuations are expected to be consistent across tasks [41, 45]. We partitioned the measured spiking activity in two components: task-tuned activity and spontaneous activity or residuals. For each channel in each session in each task, we computed task-tuned activity by replacing the firing rate of each trial with the mean firing rate across trials of the same experimental condition. This corresponds to a tuning profile where the trial-averaged spike rate for each condition is repeated as many times as the number of trials for that condition. We computed residuals by subtracting out the task-tuned activity from the measured spiking activity. We then computed topographies for task-tuned activity and residuals as described in the previous section, but replaced tuning profiles by task-tuned or residual activity vectors. Results reported in Fig. 3 are based on the task-tuned and residual topographies, allowing for a direct comparison between the two.

4.4.4. Assessing the consistency of functional topographies across time and across tasks

To assess whether the functional topographies are consistent across time and across tasks, we computed the Pearson correlation of tuning similarity matrices across session pairs within and between tasks (see Fig. 2g). Because the matrices are symmetric, we computed correlation coefficients using the upper triangular vector of the matrix. The consistency of functional topographies within and between tasks was estimated as the mean across session pairs. To control for array shifting across time, we only included sessions spaced apart no more than 20 days for within- as well as between-task comparisons. Given that data for some tasks were acquired more than 20 days apart (see Supplementary Fig. 2), between-task comparisons are based on two out of three tasks for each monkey.

To determine whether the estimated consistencies are significantly higher than chance, we permuted channel locations on the array to simulate the null hypothesis of no consistent spatial organization. We performed 1,000 permutations, each yielding an estimate of our test statistic under the null hypothesis. If the actual consistency fell within the top 5 percent of the simulated null distribution, we rejected the null hypothesis of no consistent spatial organization. This corresponds to a one-sided test using $p < 0.05$ for significance thresholding. For between-task consistencies, we did not only test if they were higher than expected for no consistent spatial organization, we also tested if they were lower than expected for a fully consistent spatial organization. We did so by estimating a noise ceiling for between-task consistencies. For each task pair, the noise ceiling is defined as the geometric mean of the within-task consistencies, which reflects the expected maximum between-task consistency given the noise in the data [66]. We compared the observed between-task consistencies against the noise ceiling using a one-sided one-sample t-test across session pairs. We compared consistencies of task-tuned and residual topographies using a two-sided paired-samples t-test across session pairs.

4.4.5. Assessing the spatial scale of functional topographies

To assess the spatial scale of the functional topographies, we computed spatial autocorrelation functions (ACFs) of channel tuning profiles on the array. The spatial ACFs were computed in a cross-validated fashion to reduce the impact of spatially autocorrelated noise. Trials in a measurement session were split into two halves. Tuning profiles were estimated for each half. We used global Moran's I as a measure of spatial autocorrelation [67], which is defined as:

$$I = \frac{N}{\sum_i \sum_j w_{ij}} \frac{\sum_i \sum_j w_{ij} (X_{Ai} - \bar{X_A})'(X_{Bj} - \bar{X_B})}{\sum_i (X_{Ai} - \bar{X_A})'(X_{Bi} - \bar{X_B})}$$

where N is the total number of channels on the array; X_{Ai} is the mean-centered tuning profile for the ith channel in the first half; X_{Bj} is the mean-centered tuning profile for the jth

channel in the second half; $\bar{X}$ is the mean-centered tuning profile averaged across all channels, estimated for each half separately; and w_{ij} is either 1 or 0 (if channel i and channel j are spaced within a specified spatial range, w_{ij} equals 1; otherwise 0).

The numerator estimates the cross-validated covariances of tuning profiles between channel pairs spaced at a certain distance. The denominator estimates the cross-validated variances of tuning profiles across all the channels on the array. To account for the fact that some channels appear more frequently than others in the covariance estimates, the variance can be estimated using the following formula as in [49]:

$$\frac{\sqrt{\sum_i \sum_j w_{ij}(X_{Ai} - \bar{X_A})'(X_{Bi} - \bar{X_B}) \sum_i \sum_j w_{ij}(X_{Aj} - \bar{X_A})'(X_{Bj} - \bar{X_B})}}{\sum_i \sum_j w_{ij}}$$

Using this formula for estimating the variance ensures that I is bounded between -1 and 1. A positive I indicates that channels spaced at the given distance have similar tuning profiles; a negative I indicates that channels spaced at the given distance have dissimilar tuning profiles.

To compute spatial ACFs, we systematically varied the spatial distance between channels, from including only immediately neighbouring channels ($0 <$ distance ≤ 0.4 mm) to channels that are spaced apart more than one channel width but no more than two (0.4 mm $<$ distance ≤ 0.8 mm), continuing these steps up and till nine channel widths (3.2 mm $<$ distance ≤ 3.6 mm), and computed the spatial autocorrelation for each distance. At distance 0, the cross-validated covariances are the same as cross-validated variances, leading to a spatial autocorrelation of 1. We computed an ACF for each array in each session in each task.

To quantify the spatial scale of the functional topographies, we fit a Laplacian function to the spatial ACFs. The Laplacian function captures the exponential decay of channel tuning similarity as the distance between channels increases. The Laplacian function used is defined as follows:

$$f(d) = 1.02 \times e^{-\frac{d}{s}} - 0.02$$

where $d \geq 0$ reflects the distances; and s is a fitted value, reflecting the smoothness of the curve. The full-width-at-half-maximum (FWHM) of the Laplacian curve can be computed using s via the following formula:

$$FWHM = 2 \times s \times \ln(2)$$

The FWHM of a fitted Laplacian function is equivalent to the FWHM of the 2-dimensional kernel required to smooth an array whose channel tuning profiles are spatially independent, to yield the degree of spatial autocorrelation we observe in the data [50]. We therefore use the FWHM as an estimator of spatial scale.

4.4.6. Mapping tuning similarity on the array

To visualize the spatial structure of tuning similarity on the arrays, we converted the correlations in the tuning similarity matrix to correlation distances, and applied 2D multidimensional scaling (MDS) to the distances. Channels were colour-coded based on their location in the 2D MDS space and projected back to the arrays. In the colour space, hue reflects polar angle, and saturation reflects eccentricity. Similar colours indicate similar tuning profiles. As a check, we computed the variance explained in the high-dimensional distances by the low-dimensional distances for a range of MDS dimensions (1-10). We computed the variance explained by first correlating the correlation distances between points (channels) in the original high-dimensional space with the Euclidean distances between points in the low-dimensional MDS space, and then squaring the correlation coefficients [41]. The 2D MDS space explains around 80% of the variance in the original high-dimensional space across monkeys, arrays, and tasks.

4.4.7. Linking the functional topographies to structural maps

To relate the spatial scale of the observed functional topographies to prior anatomical tracing work, we repeated the spatial autocorrelation analysis on a reconstructed structural map of afferent input to macaque LPFC [42]. This structural map shows the terminal field distributions of callosal fibers projecting from the principal sulcus in one hemisphere to the principal sulcus in the other hemisphere [51]. The fibers terminate in a stripe-like pattern, reflecting interdigitation of the contralateral callosal fibers with associational fibers from the ipsilateral parietal cortex [43]. The structural map suggests the existence of cortical columns in LPFC, which have been reported to have a width of 300 to 700 microns [42, 51, 44, 43, 68]. To assess the spatial scale of the structural map, we randomly sampled 210 cortical patches with 4×4 mm^2 coverage from the map, simulating array placements. Sampled cortical patches were downsampled to 10×10 to match the measurement resolution of the Utah arrays used in our study. We assessed the spatial ACF for each sampled patch and plotted the distribution of ACFs across samples. We then examined whether the ACFs observed for the task-tuned topographies fall within the distribution derived from the anatomical tracing map (see Fig. 5).

6. Data availability

Source Data are provided with this paper. Raw data supporting the findings of this paper are available from MR, BC, RL, JMT upon reasonable requests.

7. Code availability

MATLAB code used for analyzing the data is available from github (https://github.com/jkderrick028/topoPFC) or from JDX.

8. Acknowledgements

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canadian Institutes of Health Research (CIHR), Neuronex (ref. FL6GV84CKN57) grants, Ontario Graduate Scholarship, Jonathan & Joshua Memorial Graduate Scholarship in Mental Health Research, Mitacs Globalink Graduate Fellowship and McGill David G. Guthrie Fellowship. This work was further supported by BrainsCAN at the University of Western Ontario through the Canada First Research Excellence Fund (CFREF).

9. Author contributions

JDX, TWS, JMT, MM conceptualized the work. MR designed the VWM task, collected and preprocessed the data. BC, RAG designed the CDM task, collected and preprocessed the data. RL designed the ODR task, collected and preprocessed the data. RAG performed surgical implantations. JMT contributed to experimental design. JDX analyzed the data and wrote the manuscript. MHM, LM, JD, MM provided analysis advice. JDX, TWS, JMT, MM edited the manuscript.

10. Competing interests

The authors declare no competing interests.

References

[1] D. H. Hubel, T. N. Wiesel, Receptive fields and functional architecture of monkey striate cortex, The Journal of Physiology 195 (1968) 215–243. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1113/jphysiol.1968.sp008455.

[2] M. M. Merzenich, J. F. Brugge, Representation of the cochlear partition on the superior temporal plane of the macaque monkey, Brain Research 50 (1973) 275–296.

[3] J. Duncan, An adaptive coding model of neural function in prefrontal cortex, Nature Reviews Neuroscience 2 (2001) 820–829. Number: 11 Publisher: Nature Publishing Group.

[4] D. J. Freedman, M. Riesenhuber, T. Poggio, E. K. Miller, Categorical Representation of Visual Stimuli in the Primate Prefrontal Cortex, Science 291 (2001) 312–316. Publisher: American Association for the Advancement of Science.

[5] E. K. Miller, J. D. Cohen, An Integrative Theory of Prefrontal Cortex Function, Annual Review of Neuroscience 24 (2001) 167–202. _eprint: https://doi.org/10.1146/annurev.neuro.24.1.167.

[6] T. Lennert, R. Cipriani, P. Jolicoeur, D. Cheyne, J. C. Martinez-Trujillo, Attentional Modulation of Neuromagnetic Evoked Responses in Early Human Visual Cortex and Parietal Lobe following a Rank-Order Rule, Journal of Neuroscience 31 (2011) 17622–17636. Publisher: Society for Neuroscience Section: Articles.

[7] T. Lennert, J. C. Martinez-Trujillo, Prefrontal Neurons of Opposite Spatial Preference Display Distinct Target Selection Dynamics, The Journal of Neuroscience 33 (2013) 9520–9529.

[8] P. S. Goldman-Rakic, Topography of Cognition: Parallel Distributed Networks in Primate Association Cortex, Annual Review of Neuroscience 11 (1988) 137–156. _eprint: https://doi.org/10.1146/annurev.ne.11.030188.001033.

[9] J. Fuster, The Prefrontal Cortex, Academic Press, 2015.

[10] M. Rigotti, O. Barak, M. R. Warden, X.-J. Wang, N. D. Daw, E. K. Miller, S. Fusi, The importance of mixed selectivity in complex cognitive tasks, Nature 497 (2013) 585–590. Number: 7451 Publisher: Nature Publishing Group.

[11] S. Fusi, E. K. Miller, M. Rigotti, Why neurons mix: high dimensionality for higher cognition, Current Opinion in Neurobiology 37 (2016) 66–74.

[12] B. A. Wandell, S. O. Dumoulin, A. A. Brewer, Visual Field Maps in Human Cortex, Neuron 56 (2007) 366–383.

[13] A. A. Brewer, W. A. Press, N. K. Logothetis, B. A. Wandell, Visual areas in macaque cortex measured using functional magnetic resonance imaging, The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 22 (2002) 10416–10426.

[14] C. Fang, X. Cai, H. D. Lu, Orientation anisotropies in macaque visual areas, Proceedings of the National Academy of Sciences 119 (2022) e2113407119. Publisher: Proceedings of the National Academy of Sciences.

[15] A. H. Bell, F. Hadj-Bouziane, J. B. Frihauf, R. B. H. Tootell, L. G. Ungerleider, Object Representations in the Temporal Cortex of Monkeys and Humans as Revealed by Functional Magnetic Resonance Imaging, Journal of Neurophysiology 101 (2009) 688–700.

[16] P. Bao, L. She, M. McGill, D. Y. Tsao, A map of object space in primate inferotemporal cortex, Nature 583 (2020) 103–108. Number: 7814 Publisher: Nature Publishing Group.

[17] D. J. Margolis, H. Lütcke, K. Schulz, F. Haiss, B. Weber, S. Kügler, M. T. Hasan, F. Helmchen, Reorganization of cortical population activity imaged throughout long-term sensory deprivation, Nature Neuroscience 15 (2012) 1539–1546.

[18] D. B. T. McMahon, I. V. Bondar, O. A. T. Afuwape, D. C. Ide, D. A. Leopold, One month in the life of a neuron: longitudinal single-unit electrophysiology in the monkey visual system, Journal of Neurophysiology 112 (2014) 1748–1762. Publisher: American Physiological Society.

[19] L. Cossell, M. F. Iacaruso, D. R. Muir, R. Houlton, E. N. Sader, H. Ko, S. B. Hofer, T. D. Mrsic-Flogel, Functional organization of excitatory synaptic strength in primary visual cortex, Nature 518 (2015) 399–403. Publisher: Nature Publishing Group.

[20] K. Tanaka, Inferotemporal cortex and object vision, Annual Review of Neuroscience 19 (1996) 109–139.

[21] E. Yacoub, N. Harel, K. Uğurbil, High-field fMRI unveils orientation columns in humans, Proceedings of the National Academy of Sciences 105 (2008) 10607–10612. Publisher: Proceedings of the National Academy of Sciences.

[22] T. Sato, G. Uchida, M. D. Lescroart, J. Kitazono, M. Okada, M. Tanifuji, Object Representation in Inferior Temporal Cortex Is Organized Hierarchically in a Mosaic-Like Structure, The Journal of Neuroscience 33 (2013) 16642.

[23] C. Constantinidis, M. N. Franowicz, P. S. Goldman-Rakic, Coding Specificity in Cortical Microcircuits: A Multiple-Electrode Analysis of Primate Prefrontal Cortex, The Journal of Neuroscience 21 (2001) 3646–3655.

[24] D. A. Markowitz, C. E. Curtis, B. Pesaran, Multiple component networks support working memory in prefrontal cortex, Proceedings of the National Academy of Sciences 112 (2015) 11084–11089. Publisher: Proceedings of the National Academy of Sciences.

[25] N. Y. Masse, J. M. Hodnefield, D. J. Freedman, Mnemonic Encoding and Cortical Organization in Parietal and Prefrontal Cortices, The Journal of Neuroscience 37 (2017) 6098–6112.

[26] K. R. Bullock, F. Pieper, A. J. Sachs, J. C. Martinez-Trujillo, Visual and presaccadic activity in area 8Ar of the macaque monkey lateral prefrontal cortex, Journal of Neurophysiology 118 (2017) 15–28. Publisher: American Physiological Society.

[27] M. L. Leavitt, F. Pieper, A. J. Sachs, J. C. Martinez-Trujillo, A Quadrantic Bias in Prefrontal Representation of Visual-Mnemonic Space, Cerebral Cortex 28 (2018) 2405–2421.

[28] L. N. Driscoll, N. L. Pettit, M. Minderer, S. N. Chettih, C. D. Harvey, Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex, Cell 170 (2017) 986–999.e16. Publisher: Elsevier.

[29] H. Muysers, H.-L. Chen, J. Hahn, S. Folschweiller, T. Sigurdsson, J.-F. Sauer, M. Bartos, A persistent prefrontal reference frame across time and task rules, Nature Communications 15 (2024) 2115. Publisher: Nature Publishing Group.

[30] G. R. Yang, M. W. Cole, K. Rajan, How to study the neural mechanisms of multiple tasks, Current Opinion in Behavioral Sciences 29 (2019) 134–143.

[31] R. Luna, M. Roussy, S. Treue, J. C. Martinez-Trujillo, Reference Frames for Spatial Working Memory in the Lateral Prefrontal Cortex of primates, Journal of Vision 19 (2019) 206.

[32] M. Roussy, R. Luna, L. Duong, B. Corrigan, R. A. Gulli, R. Nogueira, R. Moreno-Bote, A. J. Sachs, L. Palaniyappan, J. C. Martinez-Trujillo, Ketamine disrupts naturalistic coding of working memory in primate lateral prefrontal cortex networks, Molecular Psychiatry 26 (2021) 6688–6703. Number: 11 Publisher: Nature Publishing Group.

[33] B. W. Corrigan, R. A. Gulli, G. Doucet, M. Roussy, R. Luna, K. S. Pradeepan, A. J. Sachs, J. C. Martinez-Trujillo, Distinct neural codes in primate hippocampus and lateral prefrontal cortex during associative learning in virtual environments, Neuron 110 (2022) 2155–2169.e4.

[34] G. Doucet, R. A. Gulli, J. C. Martinez-Trujillo, Cross-species 3D virtual reality toolbox for visual and cognitive experiments, Journal of Neuroscience Methods 266 (2016) 84–93.

[35] N. Kriegeskorte, M. Mur, P. A. Bandettini, Representational similarity analysis - connecting the branches of systems neuroscience, Frontiers in Systems Neuroscience 2 (2008). Publisher: Frontiers.

[36] N. Kriegeskorte, M. Mur, D. A. Ruff, R. Kiani, J. Bodurka, H. Esteky, K. Tanaka, P. A. Bandettini, Matching categorical object representations in inferior temporal cortex of man and monkey, Neuron 60 (2008) 1126–1141.

[37] M. Mur, M. Meys, J. Bodurka, R. Goebel, P. A. Bandettini, N. Kriegeskorte, Human Object-Similarity Judgments Reflect and Transcend the Primate-IT Object Representation, Frontiers in Psychology 4 (2013). Publisher: Frontiers.

[38] R. Kiani, H. Esteky, K. Mirpour, K. Tanaka, Object category structure in response patterns of neuronal population in monkey inferior temporal cortex, Journal of Neurophysiology 97 (2007) 4296–4309.

[39] W. A. Freiwald, D. Y. Tsao, Functional Compartmentalization and Viewpoint Generalization Within the Macaque Face-Processing System, Science 330 (2010) 845–851. Publisher: American Association for the Advancement of Science.

[40] N. Kriegeskorte, R. A. Kievit, Representational geometry: integrating cognition, computation, and the brain, Trends in Cognitive Sciences 17 (2013) 401–412. Publisher: Elsevier.

[41] R. Kiani, C. J. Cueva, J. B. Reppas, D. Peixoto, S. I. Ryu, W. T. Newsome, Natural Grouping of Neural Responses Reveals Spatially Segregated Clusters in Prearcuate Cortex, Neuron 85 (2015) 1359–1373.

[42] P. S. Goldman-Rakic, Modular organization of prefrontal cortex, Trends in Neurosciences 7 (1984) 419–424.

[43] P. S. Goldman-Rakic, M. L. Schwartz, Interdigitation of Contralateral and Ipsilateral Columnar Projections to Frontal Association Cortex in Primates, Science 216 (1982) 755–757. Publisher: American Association for the Advancement of Science.

[44] G. R. Leichnetz, R. F. Spencer, S. G. P. Hardy, J. Astruc, The prefrontal corticotectal projection in the monkey; An anterograde and retrograde horseradish peroxidase study, Neuroscience 6 (1981) 1023–1041.

[45] M. W. Cole, D. S. Bassett, J. D. Power, T. S. Braver, S. E. Petersen, Intrinsic and Task-Evoked Network Architectures of the Human Brain, Neuron 83 (2014) 238–251.

[46] B. B. Averbeck, P. E. Latham, A. Pouget, Neural correlations, population coding and computation, Nature Reviews Neuroscience 7 (2006) 358–366. Publisher: Nature Publishing Group.

[47] M. R. Cohen, J. H. R. Maunsell, Attention improves performance primarily by reducing interneuronal correlations, Nature Neuroscience 12 (2009) 1594–1600. Publisher: Nature Publishing Group.

[48] S. A. Arbuckle, J. Weiler, E. A. Kirk, C. L. Rice, M. Schieber, J. A. Pruszynski, N. Ejaz, J. Diedrichsen, Structure of Population Activity in Primary Motor Cortex for Single Finger Flexion and Extension, Journal of Neuroscience 40 (2020) 9210–9223. Publisher: Society for Neuroscience Section: Research Articles.

[49] M. King, C. R. Hernandez-Castillo, R. A. Poldrack, R. B. Ivry, J. Diedrichsen, Functional boundaries in the human cerebellum revealed by a multi-domain task battery, Nature Neuroscience 22 (2019) 1371–1378. Number: 8 Publisher: Nature Publishing Group.

[50] J. Diedrichsen, G. R. Ridgway, K. J. Friston, T. Wiestler, Comparing the similarity and spatial structure of neural representations: A pattern-component model, NeuroImage 55 (2011) 1665–1678.

[51] P. S. Goldman, W. J. H. Nauta, An intricately patterned prefronto-caudate projection in the rhesus monkey, Journal of Comparative Neurology 171 (1977) 369–385. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.901710305.

[52] H. Nili, C. Wingfield, A. Walther, L. Su, W. Marslen-Wilson, N. Kriegeskorte, A Toolbox for Representational Similarity Analysis, PLOS Computational Biology 10 (2014) e1003553. Publisher: Public Library of Science.

[53] R. Desimone, J. Duncan, Neural mechanisms of selective visual attention, Annual Review of Neuroscience 18 (1995) 193–222.

[54] T. W. Schmitz, J. Duncan, Normalization and the Cholinergic Microcircuit: A Unified Basis for Attention, Trends in Cognitive Sciences 22 (2018) 422–437.

[55] J. Duncan, D. Chylinski, D. J. Mitchell, A. Bhandari, Complexity and compositionality in fluid intelligence, Proceedings of the National Academy of Sciences 114 (2017) 5295–5299. Publisher: Proceedings of the National Academy of Sciences.

[56] G. R. Yang, M. R. Joglekar, H. F. Song, W. T. Newsome, X.-J. Wang, Task representations in neural networks trained to perform many cognitive tasks, Nature Neuroscience 22 (2019) 297–306. Number: 2 Publisher: Nature Publishing Group.

[57] R. Xu, N. P. Bichot, A. Takahashi, R. Desimone, The cortical connectome of primate lateral prefrontal cortex, Neuron 110 (2022) 312–327.e7.

[58] J. B. Levitt, D. A. Lewis, T. Yoshioka, J. S. Lund, Topography of pyramidal neuron intrinsic connections in macaque monkey prefrontal cortex (areas 9 and 46), The Journal of Comparative Neurology 338 (1993) 360–376.

[59] G. González-Burgos, G. Barrionuevo, D. A. Lewis, Horizontal Synaptic Connections in Monkey Prefrontal Cortex: An In Vitro Electrophysiological Study, Cerebral Cortex 10 (2000) 82–92.

[60] A. F. T. Arnsten, The Neurobiology of Thought: The Groundbreaking Discoveries of Patricia Goldman-Rakic 1937-2003, Cerebral Cortex 23 (2013) 2269–2281.

[61] A. Griffa, M. Mach, J. Dedelley, D. Gutierrez-Barragan, A. Gozzi, G. Allali, J. Grandjean, D. Van De Ville, E. Amico, Evidence for increased parallel information transmission in human brain networks compared to macaques and male mice, Nature Communications 14 (2023) 8216. Publisher: Nature Publishing Group.

[62] M. W. Cole, T. Ito, D. S. Bassett, D. H. Schultz, Activity flow over resting-state networks shapes cognitive task activations, Nature Neuroscience 19 (2016) 1718–1726. Number: 12 Publisher: Nature Publishing Group.

[63] R. A. Gulli, L. R. Duong, B. W. Corrigan, G. Doucet, S. Williams, S. Fusi, J. C. Martinez-Trujillo, Context-dependent representations of objects and space in the primate hippocampus during virtual navigation, Nature Neuroscience 23 (2020) 103–112. Number: 1 Publisher: Nature Publishing Group.

[64] B. M. Yu, J. P. Cunningham, G. Santhanam, S. I. Ryu, K. V. Shenoy, M. Sahani, Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity, Journal of Neurophysiology 102 (2009) 614–635. Publisher: American Physiological Society.

[65] T. Ito, J. D. Murray, Multitask representations in the human cortex transform along a sensory-to-motor hierarchy, Nature Neuroscience 26 (2023) 306–315. Publisher: Nature Publishing Group.

[66] E. Vul, C. Harris, P. Winkielman, H. Pashler, Puzzlingly High Correlations in fMRI Studies of Emotion, Personality, and Social Cognition, Perspectives on Psychological Science 4 (2009) 274–290. Publisher: SAGE Publications Inc.

[67] P. A. P. Moran, Notes on Continuous Stochastic Phenomena, Biometrika 37 (1950) 17–23. Publisher: [Oxford University Press, Biometrika Trust].

[68] N. M. Bugbee, P. S. Goldman-Rakic, Columnar organization of corticocortical projections in squirrel and rhesus monkeys: Similarity of column width in species differing in cortical volume, Journal of Comparative Neurology 220 (1983) 355–364. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.902200309.