

Orthogonal representations of sensory evidence and possible actions in macaque dorsolateral prefrontal cortex

Eric Kenji Lee (kenjilee@bu.edu)

Dept. of Psychological and Brain Sciences, Boston University
Boston, MA 02118 USA

Tian Wang (tianw@bu.edu)

Dept. of Biomedical Engineering, Boston University
Boston, MA 02118 USA

Nicole Carr (ncarr@bu.edu)

Dept. of Biomedical Engineering, Boston University
Boston, MA 02118 USA

Vivian Moosmann (vivian2201@gmx.de)

Dept. Business and Human Psychology, Hochschule Furtwangen
Villingen-Schwenningen, BW 78054 Germany

Pierre-Olivier Boucher (pbouche1@bu.edu)

Dept. of Biomedical Engineering, Boston University
Boston, MA 02118 USA

Chandramouli Chandrasekaran (cchand1@bu.edu)

Dept. of Anatomy and Neurobiology, Boston University
Boston, MA 02118 USA

Abstract

In perceptual decision-making, possible actions are often available before or after sensory evidence for what to do is known. We investigated how neural circuits flexibly combine these pieces of information by both recording from the prefrontal cortex of a monkey and analyzing a recurrent neural network (RNN) performing the same two decision-making tasks. In both tasks, the goal was the same: discriminate the predominant color of a red-green checkerboard and touch the corresponding target. In one task, targets appeared before the checkerboard and in the other task, the order was reversed with a working memory delay. We examined the same neurons and RNN units in both tasks allowing us to jointly examine them.

The RNN predicted the following representational geometry: sensory evidence and possible actions are represented orthogonally whereas the final action choice output emerges as a combination of these two factors. We uncovered similar dynamical structures in the brain: after appearance of the first stim./targ., sensory evidence and targ. config. were encoded within orthogonal subspaces of activity. After the onset of the second stimulus, activity evolved into an aligned action choice plane. To understand how such orthogonal task variable representations might be produced, we analyzed the selectivity of these neurons. Preliminary analysis suggests neurons segregate into populations with selectivity for target configuration or sensory evidence for the first stimulus, but not both. Conversely, subspace alignment is the result of the many action-selective neurons integrating both signals and leading to choice.

Decision-making often requires separate encoding of evidence and possible actions. For instance, when navigating in the real world, we may be presented with several routes to a destination, each with different costs that we must weigh to make an optimal decision. While recent work has shown that stimulus evidence is encoded separately from possible choices and chosen action in lateral prefrontal cortex (LPFC) (Charlton & Goris, 2022; Tafazoli et al., 2024), the exact dynamical and single cellular mechanisms behind how these different task variables are encoded and how that is transformed into action is unknown. We first generated a hypothesized dynamical mechanism by training a low-rank recurrent neural network (LR-RNN; Valente, Pillow, and Oostojic (2022)). Second, we show that the dynamics *in vivo* in the dorsolateral prefrontal cortex (DLPFC) of a macaque performing the same two tasks was consistent with the representation predicted by the RNN. Lastly, we explore how single unit selectivity might produce the dynamics seen at the population level.

RNN's provide a hypothesized dynamical mechanism of decision-making

We first investigated how units in a LR-RNN represent task variables, across two tasks in which the order of stimulus

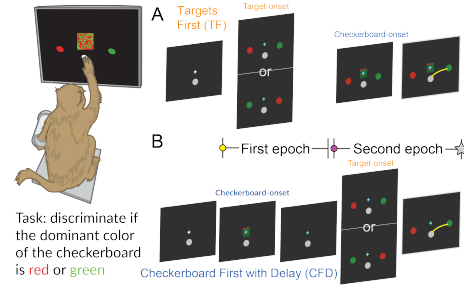


Figure 1: The two tasks: Targets First (TF; **A**) and Checkerboard First with Delay (CFD; **B**).

(stim.) and targets (targ.) are switched. In the Targets First (TF) task, red and green targ. appear with random configuration and then a checkerboard stim. appears with variable numbers of red and green squares (**Fig. 1A**). In the Check. First with Delay (CFD) task, the order is reversed and stim. is followed by a delay where the stim. disappears (**Fig. 1B**).

The RNN was given separate targ. config. and stim. signals and rank-4 constrained (**Fig. 2A**). Principal components analysis suggested the following representation: after stim. or targ. onset (yellow dot, **Fig. 2B, top**), activity spread in orthogonal directions encoding either stim. color (in CFD) or targ. config. (in TF **Fig. 2B, bot.**). After the onset of the second signal (purple dots, **Fig. 2C, top**), activity was again pushed along stim. color/targ. config. directions until each trajectory arrived at one of four locations encoding stimulus color plus reach direction demarcated by a star (**Fig. 2C, bot.**).




Figure 2: **(A)** A rank-4 LR-RNN is trained on both TF and CFD. **(B, top)** Trial-averaged RNN unit activities are projected onto the top-3 PC's, sorted according to trial outcome, and aligned to the first epoch. **(B, bottom)** These population activities are projected onto the PC1-PC2 plane. **(C)** Same as (B) but now aligned to the second trial epoch. Stars indicate end points of the trial.

Neurons in DLPFC recapitulate dynamical predictions from RNN's

To test the dynamical mechanism predicted by the RNN, we recorded from macaque DLPFC with Neuropixels probes (n=6 sessions). We used a conservative approach to select 329 well-isolated single units. We then aligned to the first or second epoch and examined the trial-averaged firing rate of each neuron for all trial outcomes (red left target choice, green left, red right, and green right). A neuron was target configuration selective (targ. selective) if it significantly differentiated ($p < 0.001$) the two target configurations; a neuron was stimulus evidence selective (stim. selective) if it significantly differentiated checkerboards with different predominant colors; and

finally a neuron was said to be action choice selective if it differentiated the direction (left vs. right) that the animal reached.

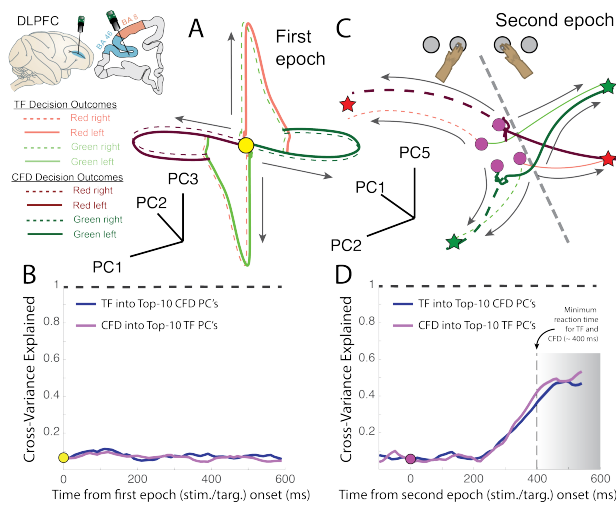


Figure 3: (A) Aligned to the first epoch (yellow dot), neural activity evolved over time in the space of first three principal components. (B) The subspace “alignment index” for TF activity in the top-10 CFD PC’s and vice versa over the same time period as A. (C) Neural activity aligned to the second epoch onset evolves toward four discrete positions in state space grouped by chosen color plus reach direction (colored stars). (D) The degree of subspace alignment increased during the second epoch until just before minimum reach reaction times.

We found that neural activity followed the hypothesized dynamics posited by the RNN. After the onset of the first epoch, neural activity spread along two orthogonal directions encoding either stim. evidence or targ. config. (compare Fig. 3A to Fig. 2B, bot.). Consistent with this visualization, we found that the subspaces encoding each task variable during this time were highly non-overlapping (Fig. 3B). We evaluated this via an “alignment index” (Elsayed, Lara, Kaufman, Churchland, & Cunningham, 2016) between the tasks calculated by projecting the activity of one task into the top-10 PC’s during the other and calculating cross-variance explained. In contrast, during the second epoch, activity evolved over time until they grouped by chosen color plus reach direction at four discrete positions in state space (colored stars, Fig. 3C) similar to the grouping seen in Fig. 2C, bot. Over this second epoch, subspaces started initially misaligned but gradually aligned contemporaneous with alignment of neural activity by chosen color and action (Fig. 3D). This alignment leveled-off just before the minimum reaction time. After the minimum reaction times, activity can be divided into left/right reaches and this difference increases over time (dashed gray line in Fig. 3C).

Single neurons fall into separate populations encoding stim. or targ. but rarely both

In order to explore potential single neuron mechanisms that produce orthogonal representation of task variables, we ex-

amined the single unit selectivity for stim., targ. config., and action choice. During the first epoch, neurons exhibited strong encoding of stim. color and targ. config. during the appearance of the checkerboard or targets respectively (Fig. 4A). We assessed the strength of encoding for either stim. color or targ. config. (or both), using a linear regression. This regression was applied in 100-ms time bins with the maximum R^2 over all trial bins shown in Fig. 4B. Only assessing units that significantly ($p < 0.001$) encoded a task variable in the first epoch, we found that most neurons significantly encoded either stim. color (blue markers) or targ. config. (orange markers) but rarely both (green markers). These results are preliminary evidence that the subspace orthogonalization found at the population level was due to a functional segregation of neurons in terms of task variable selectivity.

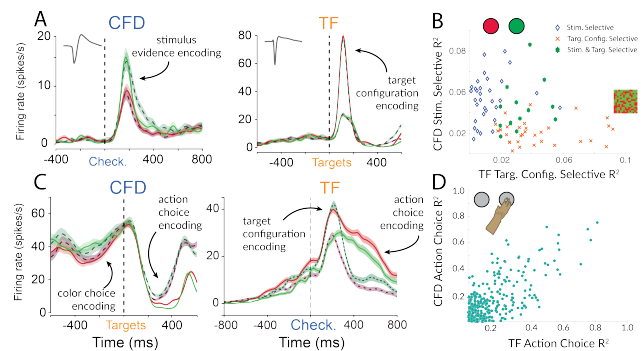


Figure 4: (A) Two exemplar DLPFC neurons with peri-stimulus time histogram in the first epoch of CFD (left) and TF (right). (B) R^2 of linear regression on single unit firing rate differences for targ. config. (x-axis) and stim. (y-axis). (C) Two exemplar DLPFC neurons with peri-stimulus time histogram in the second epoch of CFD (left) and TF (right). (D) R^2 of linear regression on single unit firing rate differences for action choice in TF (x-axis) and action choice in CFD (y-axis).

However, this previous result does not explain how the activity subspaces progressively align in the second epoch as neural trajectories converge. We next investigated if action choice selective neurons (Fig. 4C) were behind this convergence. Adopting the same regression analysis as before, we found that action choice selective neurons were highly conserved between tasks (Fig. 4D). This implies that, once both stimulus and targets are available, action choice selective neurons become active and dynamics unify to jointly signal color choice and reach direction and eventually action choice.

Conclusion: We used a novel dual task paradigm to isolate how sensory evidence and possible actions are encoded. We used a RNN model to propose a mechanism and validated it in-vivo. Our main finding was that abstract task variables are encoded in orthogonal subspaces and this is likely due to by segregated single neuron selectivity. This orthogonality gave way to subspace alignment when both pieces of information were available. Our results suggest a dynamical mechanism behind how PFC is able to encode and combined diverse types of information and use it to guide action choice.

Acknowledgments

EKL is supported by an NIH NINDS F31 award NS131018. CC is supported by a NIH/NINDS R00 award R00NS092972, NIH/NINDS R01 NS121409, NIH/NINDS NS122969, the Moorman-Simon Interdisciplinary Career Development Professorship from Boston University, the Whitehall foundation, and the Young Investigator Award from the Brain and Behavior Research Foundation.

References

- Charlton, J. A., & Goris, R. (2022). Abstract deliberation by visuomotor neurons in prefrontal cortex. doi: 10.1101/2022.12.06.519340
- Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M., & Cunningham, J. P. (2016). Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications*, 7(1), 13239. doi: 10.1038/ncomms13239
- Tafazoli, S., Bouchacourt, F. M., Ardalan, A., Markov, N. T., Uchimura, M., Mattar, M. G., ... Buschman, T. J. (2024). Building compositional tasks with shared neural subspaces. *bioRxiv*, 2024.01.31.578263. doi: 10.1101/2024.01.31.578263
- Valente, A., Pillow, J. W., & Ostojic, S. (2022). Extracting computational mechanisms from neural data using low-rank rnns. *Advances in Neural Information Processing Systems*, 35, 24072–24086.