Task preparation is reflected in neural state space dynamics

Harrison Ritz and **Aditi Jha** and **Jonathan Pillow** and **Jonathan Cohen**
Princeton Neuroscience Institute, Princeton University
contact: hritz@princeton.edu

Abstract

How do we prepare for an upcoming task? Cognitive psychologists have modeled task preparation using dynamical systems theory, however the neural correlates of these cognitive dynamics remain poorly understood. Bridging between cognitive theory and neural recordings, we used a statistical approach based on linear dynamical systems to analyze human EEG recordings during task switching. We found that the encoding of task information changes dynamically over the preparation epoch, supported by convergent evidence across both our novel state space modeling and traditional encoding analyses. These results provide a promising first step towards explicit process models of dynamic cognitive control.

Keywords: Cognitive Control, Task Switching, EEG, Dynamical Systems, Latent State Space Modeling

Introduction A longstanding literature in cognitive psychology has described task preparation as a dynamic process that configures information processing for the demands of an upcoming task (Rogers & Monsell, 1995). Recent behavioral modeling work has used dynamical systems theory to formalize preparation as movement through a cognitive state space toward an task-appropriate configuration (Musslick, Jang, Shvartsman, Shenhav, & Cohen, 2018; Steyvers, Hawkins, Karayanidis, & Brown, 2019; Jaffe, Poldrack, Schafer, & Bissett, 2023). These theories make clear predictions that task configuration dynamics will be implemented by trajectories through a neural state space, however there has been little work to date that explicitly tests this hypothesis (though see Ueltzhöffer, Armbruster-Genç, & Fiebach, 2015).

To test these the dynamic configuration hypothesis, we re-analyzed a recent task-switching EEG experiment (Hall-McMaster, Muhle-Karbe, Myers, & Stokes, 2019). We developed a novel application of latent state space modeling to formalize participants' EEG dynamics, providing a first step towards rich process models of proactive control.

Task and Sample A complete description of the sample and task are available in the original publication (Hall-McMaster et al., 2019). Briefly, 30 human participants performed a cued task-switching experiment while their brain activity was recorded using 61-channel scalp EEG. Participants responded to a compound stimulus on the basis of either its color (yellow or blue) or its shape (circle or square). Before each stimulus, participants were cued as to whether the color or shape was task-relevant ('task cue'). Note that this experiment included a reward condition that we do not analyze here.

Participants performed 10 blocks of 65 trials. We discarded any trials preceding or following an error, or that were rejected due to artifacts (leaving 469 trials on average). We used the preprocessed EEG data from the original experiment, and baselined electrodes by regressing out their average voltage during each trial's ITI (250ms - 50ms before trial onset).

Preparatory task encoding is dynamic We first validated and extended the authors' original analyses. We used Encoding Geometry Analysis (Ritz & Shenhav, 2022) to examine how task information was encoded during the cue period, testing the cross-validated reliability of multivariate encoding profiles. We regressed a design matrix with task identity (color vs shape) and cue identity (with two cue stimuli per task) on the 61-channel voltages at each timepoint, producing a time-series of encoding profiles (channel-wise patterns of regression weights). We then tested when these patterns were reliable by correlating encoding profiles across even and odd blocks (a powerful alternative to traditional encoding validation, Ritz & Shenhav, 2022).

Replicating the original experiment, we found that task information was robustly encoded during the task cue period (Fig 1A), consistent with proactive implementation of a task set (Siegel, Buschman, & Miller, 2015). We next measured how stably task information was encoded by testing the similarity of task-encoding profiles across time (i.e., temporal generalization, King & Dehaene, 2014). We correlated encoding profiles across timepoints (and blocks), finding that task encoding changed dynamically over the course of the cue period (Fig 1B; i.e., had a largely diagonal). Together, these analyses show that there is robust encoding of task information, and that task representations dynamically change over time. We next sought to characterize task dynamics with a generative state space model.

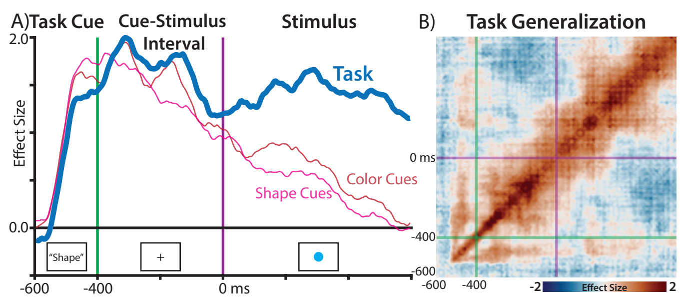

Figure 1: Task encoding. A) Task identity was reliably encoded. Y-axis reflect group-level effect size on cross-validated encoding reliability (Cohen's d), x-axis reflect time. Task cues were abstract shapes. Encoding traces are smoothed for visualization. B) Task encoding poorly generalized across time. Axes reflect same time period as panel A, color reflects the group-level effect size on cross-validated encoding alignment.

State space modeling of neural activity State space models are a statistical model of neural population activity that has growing popularity in computational neuroscience (Smith & Brown, 2003; Macke et al., 2011; Linderman, Nichols, Blei, Zimmer, & Paninski, 2019). They model time series data as arising from a set of unobserved variables, and allow us to learn the evolution of these variables over time. Here, we model participants' EEG data using a linear dynamical system (LDS) model of latent neural trajectories.

An LDS describes $\mathbf{y}_t$, the vector of neural activity at time t, as arising from the linear projection of a D-dimensional latent state $\mathbf{x}_t$. Formally, $\mathbf{y}_t = C\mathbf{x}_t + \mathbf{v}_t$ where C is a matrix of projection weights and $\mathbf{v}_t \sim \mathcal{N}(0, R)$ is Gaussian noise. The latent state $\mathbf{x}_t$ itself evolves linearly over time according to a dynamics matrix A and driven by inputs $\mathbf{u}_t$, as follows: $\mathbf{x}_t = A\mathbf{x}_{t-1} + Bu_{t-1} + \mathbf{w}_t$, where $\mathbf{w}_t \sim \mathcal{N}(0, Q)$ is Gaussian noise. $\{A, B, C, Q, R\}$ are the parameters of the LDS, describing the latent states dynamics, effect of inputs, projection of the latent states to the observed space, noise covariance of the latent states, and the observation noise, respectively.

We fit LDS models to each of our 30 participants using a custom analysis pipeline in Julia (partially ported from the SSM toolbox; Linderman, Antin, & Zoltowski, 2020). We estimated the LDS parameters using expectation-maximization, which in this case has an efficient analytic solution. The dimensionality of the latent state was selected by fitting LDS models at varying number of dimensions and selecting the model that maximizes the likelihood of a test set. We validated that our analysis pipeline could accurately recover latent trajectories in simulated data.

Neural dynamics of task encoding We estimated a series of input vectors in 10 equally spaced time bins across the Cue-Stimulus Interval (excluding the first and last 50ms). Each time bin included both a condition-independent input and task-dependent input. We report the model performance over four levels of latent dimensionality (spaced between 1 and 45 dimensions), but these trends were apparent at intermediate levels.

We found that the model's cross-validated fit increased with dimensionality (Fig 2A; confirmed with SSM toolbox). Overall, the LDS model was very accurate at reproducing single-trial EEG timeseries (Fig 2B), using the standard approach of Kalman smoothing with the estimated parameters. These predictions improved with more training data and, in simulation, with better fit to ground-truth parameters. We found that while task-independent inputs were very stable over the epoch, task-encoding dimensions changed over the course of the epoch (Fig 2C). After initial separation (likely due to differences in initial conditions), task discriminability remained relatively constant over the course of the epoch (Fig 2D). This task-specific encoding was apparent in dimensionality-reduced latent trajectories (Fig 2E).

Conclusions These exploratory analyses support the dynamic account of task preparation proposed by cognitive theories. While cognitive theories typically describe homogeneous dynamics towards a fixed location in cognitive space, we find that task subspaces are transformed over the preparation epoch. Future work should compare different models of task dynamics, and leverage this approach for spatially-resolved data (e.g., MEG). Together, the dynamical systems analyses deployed here are a first step towards observing the kind of neural dynamics that to date have only been measured using invasive recordings, opening the opportunity for richer process model of how people align neural information processing with cognitive task demands.

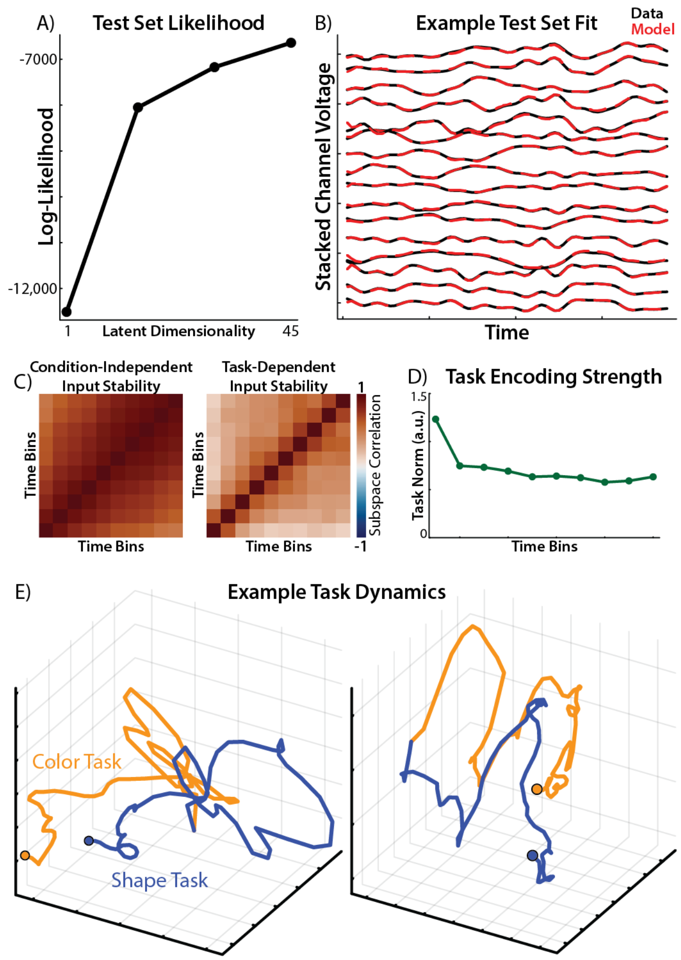

Figure 2: State Space Dynamics. A) Test-set likelihood across latent dimensionality. B) Example test-set trial timeseries, showing a subset of electrodes. Black reflects observed voltage timeseries, red reflects model predictions. Traces are offset along the y axis. C) Correlation of B across latent modes, reflecting alignment of condition-independent (left) and task-dependent (right) enocoding profiles across time bins. D) Standard deviation of B over latent modes, reflecting the magnitude of task encoding. E) Estimated latent trajectories ($\mathbf{x}_{1:T}$) for two example participants, averaged within-task and dimensionality-reduced using SVD. Circles indicate endpoints.

Acknowledgments

Thanks to Sam Hall-McMaster for making his project open source and providing helpful guidance. Thanks to Romy Frömer for advice on baselining. HR is supported by CV Starr Fellowship. This project is dedicated to the memory of Mark Stokes.

References

Hall-McMaster, S., Muhle-Karbe, P. S., Myers, N. E., & Stokes, M. G. (2019, October). Reward boosts neural coding of task rules to optimize cognitive flexibility. *J. Neurosci.*, *39*(43), 8549–8561.

Jaffe, P. I., Poldrack, R. A., Schafer, R. J., & Bissett, P. G. (2023, January). Modelling human behaviour in cognitive tasks with latent dynamical systems. *Nature Human Behaviour*, 1–15.

King, J.-R., & Dehaene, S. (2014, April). Characterizing the dynamics of mental representations: the temporal generalization method. *Trends Cogn. Sci.*, *18*(4), 203–210.

Linderman, S., Antin, B., & Zoltowski, D. (2020). *SSM: Bayesian learning and inference for state space models.* Retrieved from `http://github.com/slinderman/ssm`

Linderman, S., Nichols, A., Blei, D., Zimmer, M., & Paninski, L. (2019). Hierarchical recurrent state space models reveal discrete and continuous dynamics of neural activity in c. elegans. *BioRxiv*.

Macke, J. H., Buesing, L., Cunningham, J. P., Yu, B. M., Shenoy, K. V., & Sahani, M. (2011). Empirical models of spiking in neural populations. *Adv. Neural Inf. Process. Syst.*, *24*.

Musslick, S., Jang, S. J., Shvartsman, M., Shenhav, A., & Cohen, J. D. (2018). Constraints associated with cognitive control and the stability-flexibility dilemma. In *CogSci.* shenhavlab.org.

Ritz, H., & Shenhav, A. (2022, December). *Orthogonal neural encoding of targets and distractors supports multivariate cognitive control.*

Rogers, R. D., & Monsell, S. (1995). Costs of a predictible switch between simple cognitive tasks. *J. Exp. Psychol. Gen.*, *124*(2), 207.

Siegel, M., Buschman, T. J., & Miller, E. K. (2015, June). Cortical information flow during flexible sensorimotor decisions. *Science*, *348*(6241), 1352–1355.

Smith, A. C., & Brown, E. N. (2003, May). Estimating a state-space model from point process observations. *Neural Comput.*, *15*(5), 965–991.

Steyvers, M., Hawkins, G. E., Karayanidis, F., & Brown, S. D. (2019, August). A large-scale analysis of task switching practice effects across the lifespan. *Proc. Natl. Acad. Sci. U. S. A.*.

Ueltzhöffer, K., Armbruster-Genç, D. J. N., & Fiebach, C. J. (2015, June). Stochastic dynamics underlying cognitive stability and flexibility. *PLoS Comput. Biol.*, *11*(6), e1004331.