Neural population dynamics in prefrontal cortex and hippocampus during paired-associate learning

Yue Liu[1,2], Scott L. Brincat[4,5], Earl K. Miller[4,5], and Michael E. Hasselmo[2,3]

[1] *Department of Physics, Boston University, Boston, MA 02215*
[2] *Center for Systems Neuroscience, Boston University, Boston, MA 02215*
[3] *Department of Psychological and Brain Sciences, Boston University, Boston, MA 02215*
[4] *The Picower Institute for Learning and Memory, Massachusetts Institute of Technology, Cambridge, MA 02139*
[5] *Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139*

Abstract

Large-scale neuronal recording techniques have enabled discoveries of population-level mechanisms for neural computation. However it is not clear how these mechanisms form by trial and error learning. In this paper we present an initial effort to characterize the change of population activity in monkey prefrontal cortex (PFC) and hippocampus (HPC) during the learning phase of a paired-associate task. To analyze the population data, we introduce the normalized distance, a dimensionless metric that describes the encoding of cognitive variables from the geometrical relationship among neural trajectories in state space. It is found that PFC exhibits a more sustained encoding of task-relevant variables whereas HPC only transiently encodes the identity of the stimuli. Moreover, PFC and HPC also exhibit learning-dependent changes for the encoding of some task variables. This study shows the feasibility of using normalized distance as a metric to characterize and compare population level encoding of task variables, and suggests further directions to explore the learning-dependent changes in the population activity.

1 Introduction

With the development of experimental techniques for recording the activity of a large number of neurons, researchers have started exploring the possibility that the collective dynamics of interacting populations of neurons forms basic units for some neural computations (Sussillo (2014)). The collective dynamics are usually described by neural trajectories in state space which represent firing rates evolving through time. By looking at the collective neural dynamics through the lens of dynamical systems, several studies have identified cognitive functions with familiar concepts in dynamical systems. For example, in Mante et al. (2013), it was shown that the monkey prefrontal network performed a context-dependent decision making task by forming a pair of line attractors for the two contexts (Mante et al. (2013)). In Remington et al. (2018), the authors showed that in an interval production task, monkey frontal cortex circuits encoded the information about the time interval to be reproduced in the initial condition of the neural population dynamics, and the neural dynamics for different reproduced time intervals were represented by parallel neural trajectories with different speeds (Remington et al. (2018)).

Despite great progress in revealing collective functional features in neural computation, it remains unclear how these collective features are formed during training. Although it has been shown that similar features are present in recurrent neural networks (RNNs) trained on the same task by backpropagation, the learning dynamics of these RNNs are likely to differ from real neural circuits. To elucidate the solutions to this problem, one should look at how population dynamics change during the learning phase of a task. There exists studies that look at population level changes during motor learning (Sadtler et al. (2014); Golub et al. (2018); Vyas et al. (2018)), but similar work for cognitive learning has been scarce (although see Durstewitz et al. (2010)). In this paper

we present an initial effort to characterize population level dynamics during the learning phase of a cognitive task.

The task we analyzed is a paired associate task where a monkey learned associations between a pair of randomly chosen visual stimuli and a third visual stimulus. For analysis of the neural recording data, we introduced a dimensionless metric which we called normalized distance (ND) that describes the geometric relationships among neural trajectories for different experimental conditions. Unlike decoding methods, ND characterizes the encoding of all task variables in the population based directly on the geometry of the neural code. The dimensionless property of ND also enables comparisons of information in population codes between different learning stages as well as different brain areas. Using this metric, we then compared population level dynamics between prefrontal cortex (PFC) and hippocampus (HPC) as well as across learning stages in terms of the information content in the population codes. Our results reveal a series of differences in the dynamics of the information content between PFC and HPC as well as across learning stages. It also shows that normalized distance is a robust way of measuring information content in population codes with high temporal resolution in the face of noisy neural data.

2 Methods

2.1 Task and recording

We analyzed neural recording data from Brincat and Miller (2015). In that study, two macaque monkeys were trained to perform an object paired-associate learning task that required them to learn arbitrary associations between pairs of visual images of objects. On each day, six novel images were chosen. Four of them were randomly chosen to be assigned as the cue objects and the other two were assigned as the associate objects. Two random cue objects were paired with a random associate object, forming a 4:2 mapping from cues to associates (Figure 1a). This design enables us to distinguish neural activity that encodes the identity of the cues versus activity that encodes the associate that each particular cues corresponds to. For the rest of the paper we call the latter the "category" information because it represents which category particular cues belong to.

The structure of a trial is illustrated in Figure 1b. During each trial, after a 500 ms fixation period, the monkey was first presented with one of the cue stimuli for 500 ms, and then after a delay of 750 ms was presented with one of the associate objects for another 500 ms. The monkey should indicate whether the previous two stimuli are associated with each other by making a saccade to the indicated position on the screen if they match. The monkey was rewarded with water if the response was correct (Figure 1b).

During each recording session, the monkey must learn the novel associations from trial and error. The monkey was able to learn the associations above chance in all recording sessions (Figure 1c).

During each recording session, microelectrodes were lowered into the lateral prefrontal cortex and hippocampus and recorded spike and LFP signals while the monkeys performed the paired-associate learning task. In this paper we only analyzed the spiking data.

2.2 Normalized distance

In this section we introduce the metric we used to characterize the information content in a neural population. This metric is computed from distance between neural trajectories in state space, but normalized properly to account for the true neural information about task variables. Hence we named it "normalized distance" (ND). This metric is similar to the one used to characterize the community structure in fMRI data (Schapiro et al. (2016)), as well as the "abstraction index'' used in analyzing electrophysiological data (Bernardi et al. (2018)). All of these metrics describe the community structure of neural states organized by task variables. The difference between ND and the previous two metric is that ND is computed in continuous time. Therefore it reveals the dynamics of neural information.

When analyzing recordings from a population of neurons, it is often convenient to represent the simultaneous activity of all neurons as a point in a "state space". A state space is a high dimensional space where each axis represents the activity (in our case binned spike counts) of one

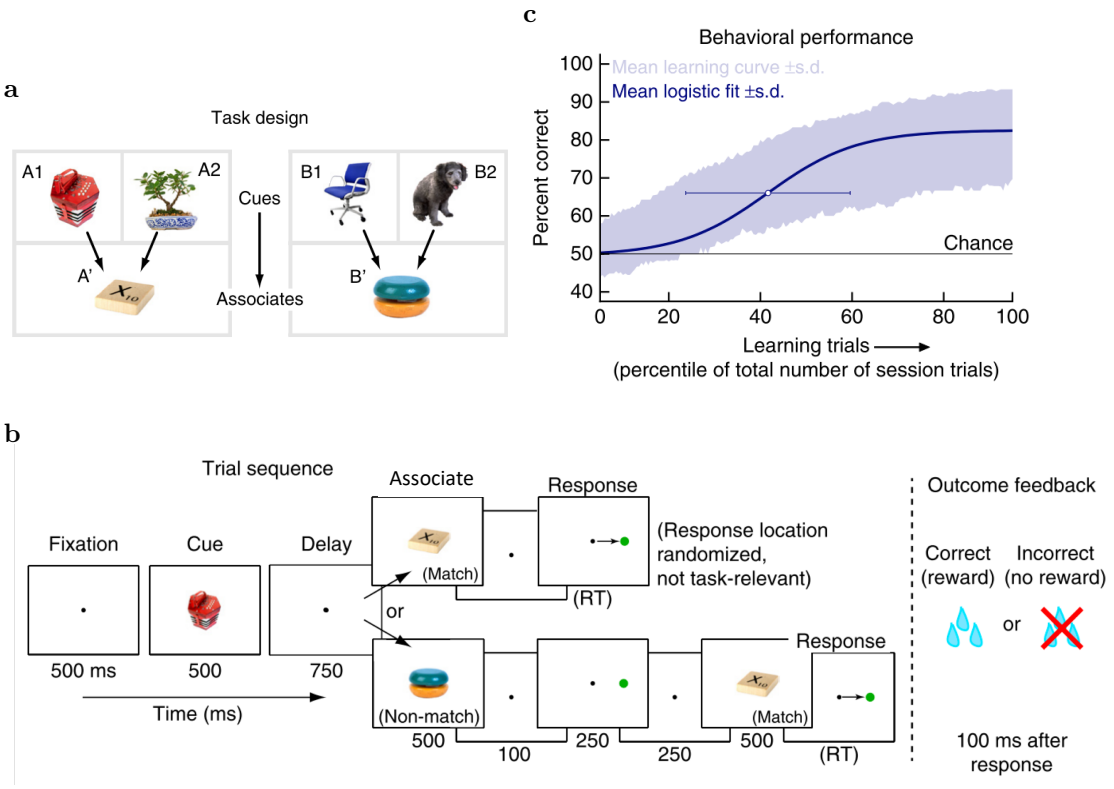

Figure 1: Overview of the task. On each session, six random stimuli are arranged into two-to-one mappings (**a**) and the monkey is required to learn the mappings via trial and error. On each trial, the monkey receives a cue and a choice stimuli separated by a 750ms delay, after which it has to make a decision of whether the two stimuli match by a saccade response for match (**b**). During each session the monkey's performance gradually increases (**c**). Adapted from Brincat et al., 2015.

neuron. Over time, the population dynamics can be represented by a trajectory through the state space.

The ND for two experimental variables X and Y is a function through time. At each point in time, it is defined to be the average Euclidean distance between pairs of neural trajectories encoding variables X and Y, normalized by the average Euclidean distance between pairs of neural trajectories that encode either X or Y (see Figure 2 for a graphical illustration).

$$ND(X,Y) = \frac{\langle D_{ij} \rangle}{\langle d_{ij} \rangle} = \frac{\left\langle |\mathbf{x_i}(t) - \mathbf{x_j}(t)| \right\rangle_{i \in X, j \in Y}}{\left\langle |\mathbf{x_i}(t) - \mathbf{x_j}(t)| \right\rangle_{i,j \in X \text{ or } i,j \in Y}}. \tag{1}$$

In the equation above, i and j index neural trajectories, D_{ij} and d_{ij} represent the Euclidean distance between neural trajectories i and j when they encode the same or different cognitive variable, respectively. $\mathbf{x_i}(t)$ and $\mathbf{x_j}(t)$ are population vectors at time t for neural trajectories i and j. Brackets denote averaging, and vertical lines denote the magnitude of a vector.

ND provides an accurate way of computing the information content in the population code directly from the geometry of the neural data. In Equation 1, the numerator is the average variability in the population code induced by the task variable of interest (here X and Y). The denominator is the average variability induced by all the other factors when the task variable of interest in fixed. The extra variability could come from nested variables within X and Y, trial history, or simply intrinsic neuronal noise. An ND(X,Y) greater than 1 indicates that the population code carries information about variables X versus Y, because there is extra variability in the population code caused by the difference between X and Y beyond the variability caused by all the other factors. On the other hand, an ND(X,Y) close to 1 indicates that the population code does not have extra variability beyond what is caused by these factors, and therefore does not carry actual information about X versus Y.

It is worth emphasizing again that the normalizing part in ND (the denominator in Equation 1) is crucial. A large Euclidean distance between two trajectories (the numerator of Equation 1) does not necessarily indicate that the neural population encodes information about the experimental condition. For example, consider the case where there are two nested variables within each experimental condition A and B (Figure 2). A large distance between pairs of neural trajectories encoding A and B does not necessarily indicate that the population carries information about A versus B. It might be that the distance between trajectories encoding each of the four trial types is equally large (Figure 2, right), in which case the population code equally discriminates between each trial type, and therefore does not carry information about A versus B. Only when the inter-variable distance is large compared to intra-variable distances does the population carry information about variable A versus B (Figure 2, left). The same argument carries over when the intra-variable distances are caused by other source of variability such as trial history or intrinsic neuronal noise.

3 Results

We applied the normalized distance metric introduced above to compute the time evolution of the information encoded in the neural population. We looked at the encoding of every cognitive variable in the task, namely cues (A_1 vs. A_2 vs. B_1 vs. B_2), associates (A' vs. B'), category information (if the first stimulus belongs to associate A or B) and decision/movement preparation (match vs. non-match).[1]

We found that PFC population encodes information about all of the cognitive variables, whereas HPC population only encodes the identity of the cue and associate stimuli. We also found the neural information in PFC is more persistent compared with HPC.

To look at learning-dependent change in neural activity, we divided the learning into 3 stages (early, mid, late) by evenly dividing all trials within a session. We also divided all trials into 15 stages to obtain a more fine-grained picture of the learning dynamics. Across learning, we found that the neural information about decision/movement preparation in PFC gradually increased,

[1] In this experiment we cannot tease apart the neural information for decision variable versus general preparation of an upcoming movement.

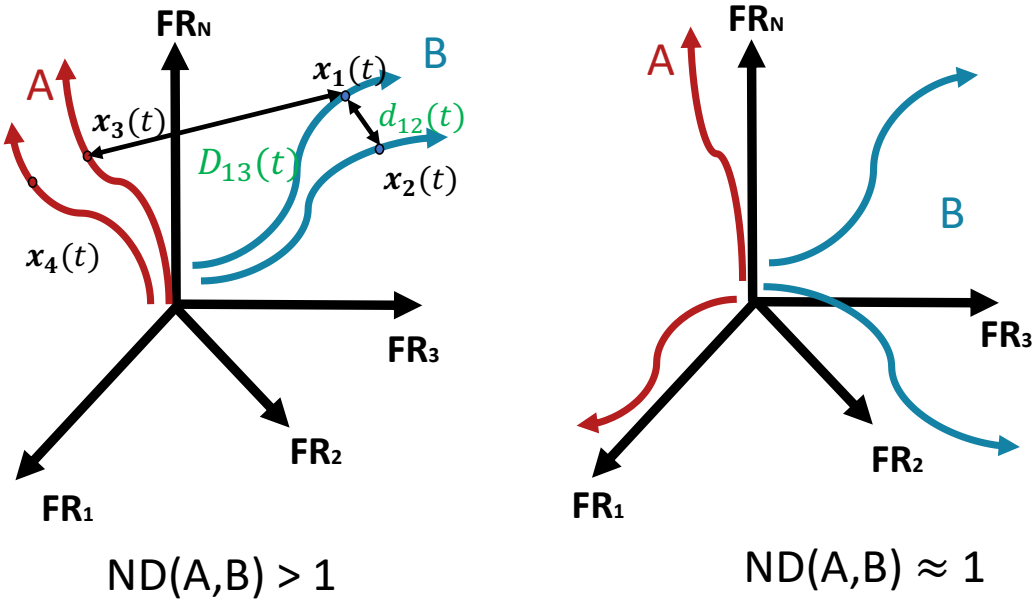

$$\mathrm{ND}(\mathrm{A},\mathrm{B}) = \frac{\langle \mathrm{D_{ij}(t)} \rangle}{\langle \mathrm{d_{ij}(t)} \rangle}$$

Figure 2: **Normalized distance (ND) characterizes the geometry of neural trajectories in state space.** Raw distance between neural trajectories is not indicative of the information content encoded by the neuronal population. A large distance in the state space between two trajectories that correspond to different experimental conditions does not indicate the population code discriminates between the two conditions. For example, if the distances between within-condition neural trajectories are comparable to that of across-condition neural trajectories (right), the population does not carry information about the task variable that is different for the two conditions (here A and B). Only when the inter-condition distances are larger than the intra-condition distances does the population encodes the information about that task variable (left).

whereas the information about the task structure seemed to be present early in the learning. In HPC, the neural information about the identity of the associate stimuli increases with learning.

Neural trajectories were obtained from condition-averaged firing rates. Therefore there are 8 neural trajectories in each session corresponding to combinations of 4 cues and 2 associates. A given $ND(X, Y)$ was computed by partitioning all 8 neural trajectories into groups that correspond to each variable, and then normalize the average across-group distance with average within-group distance, as in Equation 1. For the population analysis, only sessions where 5 or more neurons were simultaneously recorded were used in computing ND, and for a given session only correct trials were included in calculating the condition-averaged response. In the subsequent sections the time evolution of ND for each cognitive variable across learning stages and brain areas will be presented in turn.

3.1 Cue encoding

The ND among the 4 cue stimuli was computed as a function of time. The normalization factor in Equation 1 was calculated from the distance between trajectories encoding the same cue but different associate stimuli. According to the discussion above, an ND larger than 1 implies that the population contains information about the identity of the cues. On the other hand, an ND close to 1 implies that the population does not carry information about the cues.

Figure 3 shows the ND among cues as a function of time. In both PFC and HPC, ND(cues) rapidly goes up when the cue is presented at t=500ms, reflecting the encoding of the cue information at stimulus presentation. The information about the cues then gradually decays to baseline. For PFC, it rapidly decreases after the subsequent stimulus (associate) is presented. Notably, PFC has a much more sustained ND than HPC. This implies that the network timescale in PFC may be longer than HPC. There are no significant learning-dependent changes in ND for both PFC and HPC during the cue period. Overall, the results show that both PFC and HPC populations carry transient information about the cues, and that information lasts longer in PFC.

3.2 Associate encoding

The ND between the 2 associate stimuli $ND(A', B')$ was computed as in Equation 1, where the denominator is the average distance between neural trajectories encoding the same associate stimuli but different cue stimuli. As shown in Figure 4, the ND between the two associates (A' vs. B') rapidly increases when the associate is presented (t=1750ms). Similar to cue encoding, the ND in PFC shows sustained significance for a longer period of time than that in HPC. However, different from cue encoding, the ND in HPC has an interesting learning-dependent change: the average ND during the associate presentation period is significantly (Mann-Whitney U test, $p < 0.01$) smaller during early learning period (the first 1/3 of the trial) than mid (the middle 1/3) and late (the last 1/3) learning periods. By dividing the session into 15 learning stages, it is seen that the $ND(A', B')$ approximately grows linearly (Figure 4d, linear fit of the means, $p = 8 \times 10^{-5}$, Figure 4e, raw data for each session). This possibly reflects some potential constraints on HPC in encoding sequential inputs. Future research studies that vary the delay interval could elucidate the nature of this problem.

3.3 Category information encoding

We next investigated whether the information about the cue stimulus category (A1/A2 vs. B1/B2) was encoded in PFC and HPC. The ND between cue categories ($ND(A, B)$) were computed according to Equation 1. The normalizing factor in Equation 1 is the average distance between neural trajectories representing the same cue category but different cue or associate identity. As shown in Figure 5, in PFC, ND for categories increases as soon as the cue is presented (t=500ms), and is sustained throughout the delay interval. The time course is similar to ND for cues (Figure 3). Notably, the amplitude of ND is smaller for categories than for cues (compare Figure 5 with Figure 3), indicating that the neural trajectories are mainly clustered by cues, and less by categories. On the contrary, $ND(A, B)$ for HPC fluctuates around 1, indicating that HPC does not encode information about the category of the cue.

a

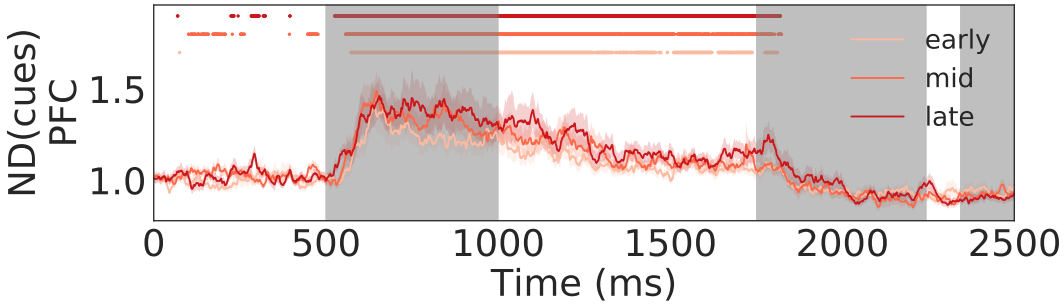

b

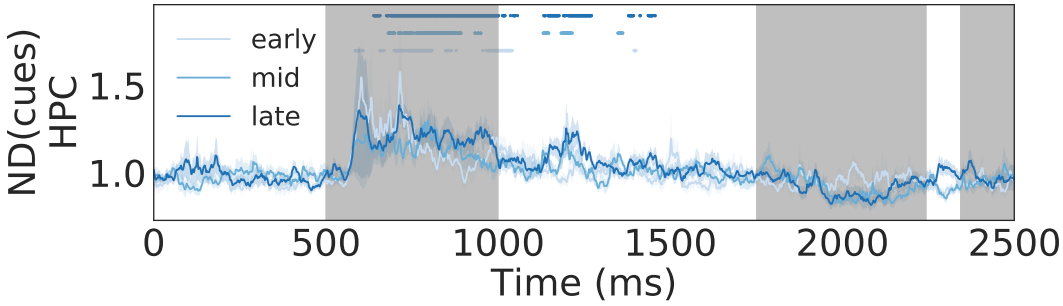

Figure 3: **Normalized distance for cues.** Both PFC (**a**) and HPC (**b**) encode cue information. The information in PFC persists longer, potentially reflecting the longer timescale of the neuronal activity in PFC compared to HPC. Shaded area shows 68% confidence interval computed from 10000 iterations of bootstrap resampling across sessions. Dots on top represent timepoints when the ND is significantly($p < 0.05$) larger than 1. Gray shaded areas indicate time periods when the cue and associate are presented and response is made.

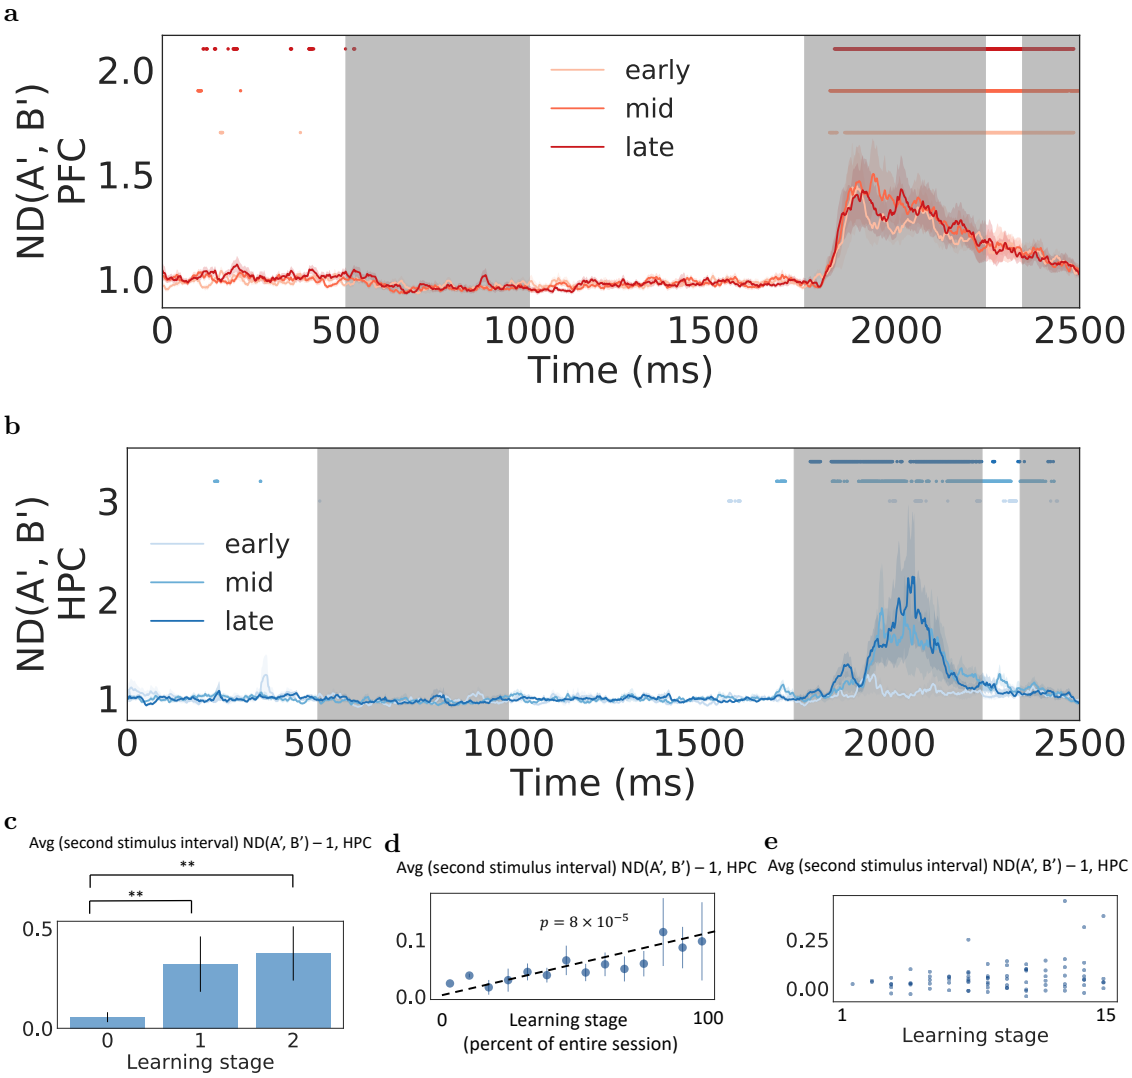

Figure 4: **Normalized distance for associates.** Both PFC (**a**) and HPC (**b**) encoded information about the identity of the associate stimulus. The information in PFC sustains longer than HPC (Shaded area shows 68% confidence interval computed from 10000 iterations of bootstrap resampling across sessions. Dots on top represent timepoints when the ND is significantly ($p < 0.05$) larger than 1. Gray shaded areas indicate time period when the cue and associate are presented and response is made). Notably, in HPC the information about the associate stimulus is stronger later in learning (**c**, ND averaged across the associate presentation period). The averaged ND grows with learning in an approximately linear fashion (**d**, linear fit for means across sessions, $p = 8 \times 10^{-5}$, **e**, raw data for each learning stage. Each data point represents a session.)

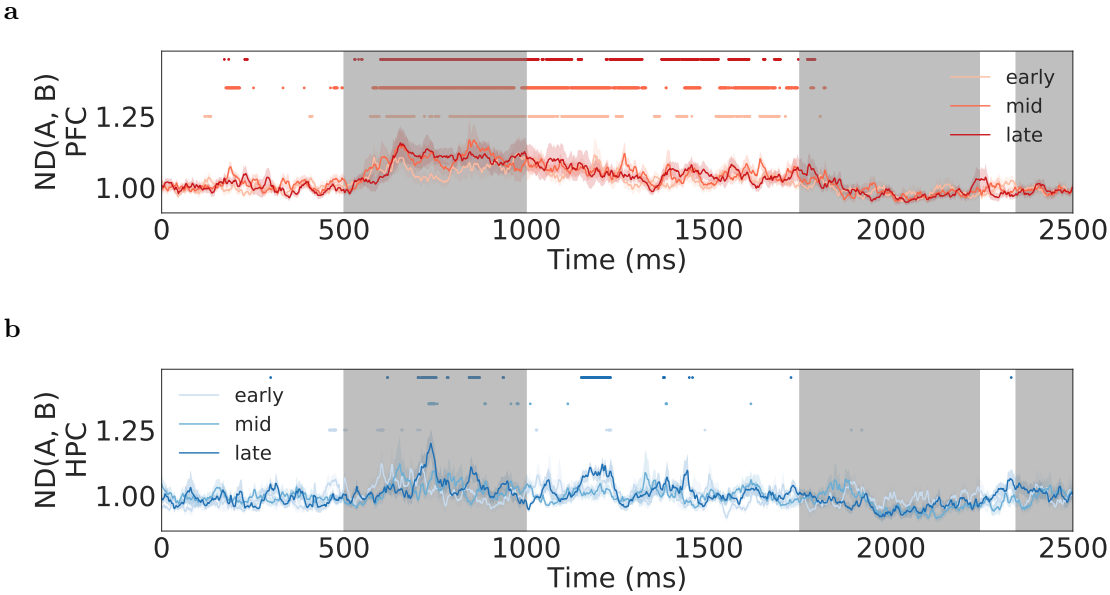

Figure 5: **Normalized distance for categories.** PFC encodes information about the category information (that A_1 and A_2 both predict A', and B_1 and B_2 predict B'), whereas HPC did not encode this information. Shaded area shows 68% confidence interval computed from 10000 iterations of bootstrap resampling across sessions. Dots on top represent timepoints when the ND is significantly($p < 0.05$) larger than 1. Gray shaded areas indicate time period when the cue and associate are presented and response is made.

The information about category reflects the learned task structure and would be expected to gradually appear with learning. Interestingly, in this dataset there seems to be no learning-dependent increase in the ND for PFC, indicating that learning about the task structure probably started very early in the training period, and the error made later were likely due to other reasons such as not paying attention or movement error.

3.4 Decision variable/movement encoding

Lastly, we looked at the ND between match and non-match trials. Since we only looked at correct trials, decision about match versus non-match is perfectly confounded with movement preparation. In this experiment there is no way to look at one of them without the influence of the other.

The ND between match and non-match trials is shown in Figure 6. In PFC, the ND between match and non-match trials starts to ramp up halfway during the second stimulus presentation period when all the information needed to form the decision is present. There is a latency of about 150 ms between the appearance of the associate information in PFC and that of the decision (compare Figure 6a with Figure 4a). This potentially indicates the time that the monkeys took to compare the incoming stimulus with the cue in the working memory. The amplitude of the ramp for the ND increases with learning (Figure 6d, e) although the effect is not significant by only comparing between 3 learning stages (Figure 6c). The increase in the amplitude of the ramp of the ND possibly reflects the increase of evidence for match versus non-match as a result of learning. In contrast, in HPC the ND between match and non-match trials fluctuates around 1 (Figure 6b). Therefore HPC does not encode the decision variable nor is responsible for movement preparation in this task.

4 Discussion

In this paper we developed a new metric called normalized distance (ND) that characterizes neural information in the population code based on the geometric organization of neural trajectories in

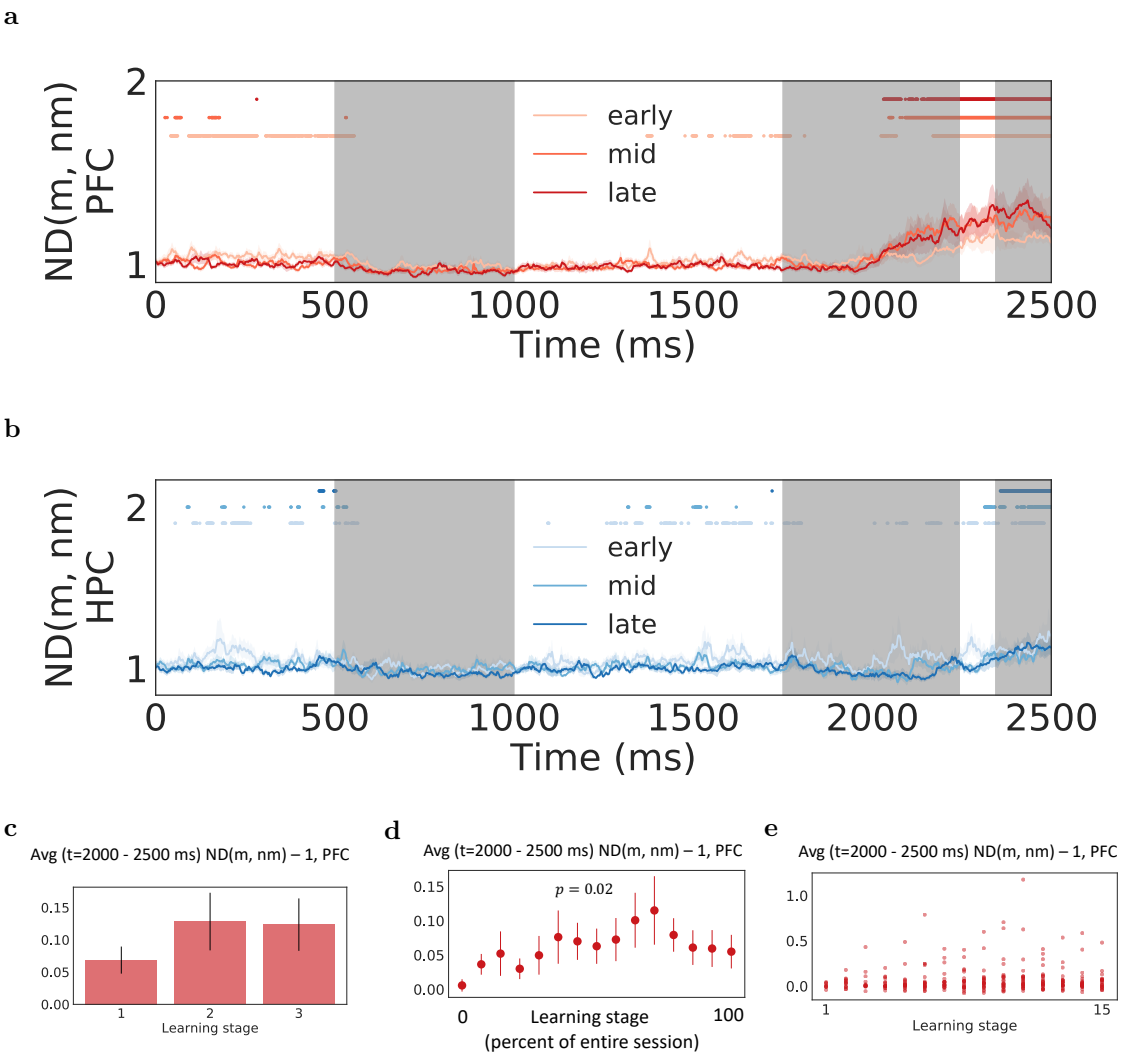

Figure 6: **Normalized distance for decision/movement.** In PFC, information about the match versus nonmatch trial type starts to ramp up halfway through the second stimulus interval (**a**). The amplitude of the ramp does not significantly differ when only comparing 3 learning stages (**c**, average ND after the presentation of the associate stimulus), but does increase with learning when dividing each session into more learning stages (**d**, linear fit for means across sessions, $p = 0.02$, **e**, raw data for each learning stage. Each data point represents a session). On the contrary, the ND for HPC fluctuates around 1, therefore HPC does not encode information about the decision/movement variable (**b**). For **a, b**: dots on top represent timepoints when the ND is significantly ($p < 0.05$) larger than 1. Gray shaded areas indicate time periods when the cue and associate are presented. Shaded area shows 68% confidence interval computed from 10000 iterations of bootstrap resampling across sessions

the state space. We then applied this metric to recording from a paired-associate task (Brincat and Miller (2015, 2016)) to compare the dynamics of the coding of different task variables across learning stages and brain areas (PFC and HPC). We found that 1) PFC exhibits a longer network timescale for object coding. 2) Object category coding is present in PFC but not HPC. 3) Coding for the associate object increases with learning in HPC but not PFC, and 4) Coding for decision/movement increases with learning in PFC but is not present in HPC.

In the original paper (Brincat and Miller (2015, 2016)), the authors used another metric, bias-corrected percentage of explained variance, to calculate neural information as a function of learning. They discovered that the object category coding was present in PFC but not HPC (Brincat and Miller (2015, 2016)), corroborating our second finding above. However, they also found that the object category coding increased with learning, which we did not find using the ND metric. It could be that the different metrics used caused different results, but it should be noted that learning effects are subtle in this experiment. Brincat and Miller also reported changes in network synchrony with learning and differences in the oscillatory modes across the two brain regions which we did not study in this work. The ND metric could in principle be applied to analyze LFP data.

We did report some results that were not reported in the original paper (results 1, 3, 4 above). In particular, the network timescale for object coding was found to be longer in PFC than HPC, potentially providing suitable temporal dynamics for PFC to integrate incoming information from other cortices. It was indeed reported that single neurons in PFC have, on average, longer time constants than motor areas (Murray et al. (2014)).

There are other studies that reported single cell activity during associative learning (Sakai and Miyashita (1991); Suzuki (2007)). These earlier studies used metrics based on single cell firing rates to correlate with learning performance and experimental conditions. In contrast, in this paper the ND metric we used characterize the distributed information in the population code.

The ND metric developed in this paper serves a similar role as decoding accuracy commonly adopted in analyzing population level data. In decoding analysis hyperplanes are constructed that separate training data from different categories as well as possible according to some objective function, and the decoding accuracy reflects how well these hyperplanes separate the held-out test data. It is noted however that in high dimensional cases where the number of neurons is comparable with the number of data points (as in our case here as well as data obtained by modern large-scale recording techniques), the decoding accuracy can be generically high (Buonomano and Maass (2009)) and thus does not necessarily reflect the underlying geometry of the neural code. It is therefore difficult to interpret results obtained by directly comparing decoding accuracies across different recording sessions. On the other hand, ND is directly computed from the geometrical configuration of the data, and therefore can be used regardless of the number of data points compared with the number of dimensions and provide a clear geometrical picture of the underlying neural code.

In this data set HPC does not encode the category information of the cues that would help predict the upcoming stimuli even after learning. This seems contradictory with the study by Stachenfeld et al. showing that HPC contains a predictive map of the environment (Stachenfeld et al. (2017)) and the study by McKenzie et al. showing rodent HPC population encodes the hierarchical structure of the task (McKenzie et al. (2014)). However, the scenario studied by Stachenfeld is a reinforcement learning task in a spatial setting where sensory experience is almost continuous. The task studied by McKenzie et al. also has a spatial component, and the population activity was observed to be largely organized by spatial context. On the other hand, the task we analyzed here is a simple sequential decision making task. It could be that in this simple setting, the HPC is not utilized to form more sophisticated relational representations (Eichenbaum (2017))

The normalized distance is directly calculated from the geometry of neuronal population responses in the state space. Therefore it provides a characterization of the degree of "tangling" of the underlying neuronal manifolds. Disentangled neuronal manifolds were argued to be crucial in forming a "good" neuronal representations for higher level processing (DiCarlo and Cox (2007)). An ND larger than 1 indicates that the neuronal manifolds representing different variables are somewhat disentangled. However, we do note that other geometric quantities such as curvature (Bernardi et al. (2018); Chaudhuri et al. (2019)) are needed to provide a complete characterization of the neuronal manifold.

Acknowledgements

This work was supported by ONR MURI award N00014-16-1-2832 (Y.L., M.E.H., E.K.M.), ONR DURIP award N00014-17-1-2304 (Y.L., M.E.H., E.K.M.) and NIMH R37MH087027 (S.L.B., E.K.M.).

References

Bernardi, S., Benna, M. K., Rigotti, M., Munuera, J., Fusi, S., Daniel, &. C., †1, S., and Zuckerman, M. B. (2018). The geometry of abstraction in hippocampus and pre-frontal cortex.

Brincat, S. L. and Miller, E. K. (2015). Frequency-specific hippocampal-prefrontal interactions during associative learning. *Nature Neuroscience*, 18(4):576–81.

Brincat, S. L. and Miller, E. K. (2016). Prefrontal Cortex Networks Shift from External to Internal Modes during Learning. *Journal of Neuroscience*, 36(37):9739–9754.

Buonomano, D. V. and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, 10(2):113–25.

Chaudhuri, R., Gercek, B., Pandey, B., Peyrache, A., and Fiete, I. (2019). The population dynamics of a canonical cognitive circuit. *bioRxiv*.

DiCarlo, J. J. and Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, 11(8):333–341.

Durstewitz, D., Vittoz, N. M., Floresco, S. B., and Seamans, J. K. (2010). Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron*, 66(3):438–48.

Eichenbaum, H. (2017). On the integration of space, time, and memory. *Neuron*, 95(5):1007–1018.

Golub, M. D., Sadtler, P. T., Oby, E. R., Quick, K. M., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P., Chase, S. M., and Yu, B. M. (2018). Learning by neural reassociation. *Nature Neuroscience*, 21(April):1–10.

Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84.

McKenzie, S., Frank, A. J., Kinsky, N. R., Porter, B., Rivière, P. D., and Eichenbaum, H. (2014). Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron*, 83(1):202–15.

Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., et al. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature neuroscience*, 17(12):1661–1663.

Remington, E. D., Narain, D., Hosseini, E. A., and Jazayeri, M. (2018). Flexible Sensorimotor Computations through Rapid Reconfiguration of Cortical Dynamics. *Neuron*, 98(5):1005–1019.e5.

Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., Yu, B. M., and Batista, A. P. (2014). Neural constraints on learning. *Nature*, 512(7515):423–426.

Sakai, K. and Miyashita, Y. (1991). Neural organization for the long-term memory of paired associates. *Nature*, 354(6349):152–155.

Schapiro, A. C., Turk-Browne, N. B., Norman, K. A., and Botvinick, M. M. (2016). Statistical learning of temporal community structure in the hippocampus. *Hippocampus*, 26(1):3–8.

Stachenfeld, K. L., Botvinick, M. M., and Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, 20(11):1643–1653.

Sussillo, D. (2014). Neural circuits as computational dynamical systems. *Current Opinion in Neurobiology*, 25:156–163.

Suzuki, W. (2007). Integrating Associative Learning Signals Across the Brain. *Hippocampus*, 251(January):235–251.

Vyas, S., Even-Chen, N., Stavisky, S. D., Ryu, S. I., Nuyujukian, P., and Shenoy, K. V. (2018). Neural Population Dynamics Underlying Motor Learning Transfer. *Neuron*, 97(5):1177–1186.e3.