Neurobiology of Learning and Memory 171 (2020) 107215

Contents lists available at ScienceDirect

Neurobiology of Learning and Memory

journal homepage: www.elsevier.com/locate/ynlme

Dorsomedial prefrontal cortex and hippocampus represent strategic context even while simultaneously changing representation throughout a task session

Brendan M. Hasz[a], A. David Redish[b,*]

[a] Graduate Program in Neuroscience, University of Minnesota Twin Cities, Minneapolis, MN, USA
[b] Department of Neuroscience, University of Minnesota Twin Cities, Minneapolis, MN, USA

ARTICLE INFO

Keywords:
Dorsomedial prefrontal cortex
Hippocampus
Learning
Memory
Neuronal representation
Prelimbic cortex

ABSTRACT

Dorsomedial prefrontal cortex (dmPFC) and hippocampus (HPC) are thought to play complementary roles in a spatial working memory and decision-making network, where spatial information from HPC informs representations in dmPFC, and contextual information from dmPFC biases how HPC recalls that information. We recorded simultaneously from neural ensembles in rodent dmPFC and HPC as rats performed a rule-switching task, and found that ensembles in dmPFC and HPC simultaneously encoded task contingencies and other time-varying information. While ensembles in HPC transitioned to represent new contingencies at the same time as rats updated their strategies to be consistent with the new contingency, dmPFC ensembles transitioned earlier. Neural representations of other time-varying information also changed faster in dmPFC than in HPC. Our results suggest that HPC and dmPFC represent contingencies while simultaneously representing other information which changes over time, and that this contextual information is integrated into hippocampal representations more slowly than in dmPFC.

1. Introduction

The prefrontal cortex (PFC) has long been thought to mediate executive function (Miller & Cohen, 2001; Dalley, Cardinal, & Robbins, 2004; Kesner & Churchwell, 2011) by participating in the storage and recall of contextual memories (Tronel & Sara, 2003; Euston, Gruber, & McNaughton, 2012; Preston & Eichenbaum, 2013) and maintaining that information in working memory (Ragozzino & Kesner, 1998; Delatour & Gisquest-Verrier, 1999; Cowen & McNaughton, 2007; Yoon, Okada, Jung, & Kim, 2008; Horst & Laubach, 2009; Urban, Layfield, & Griffin, 2014). Specifically, the dorsomedial prefrontal cortex (dmPFC) plays several different roles which support behavioral flexibility. dmPFC is important for conflict resolution (Haddon & Killcross, 2005; Haddon & Killcross, 2006; Marquis, Killcross, & Haddon, 2007; Dwyer, Dunn, Rhodes, & Killcross, 2010), as well as for learning task contingencies and adjusting behavioral strategies accordingly (Balleine & Dickinson, 1998; Floresco, Block, & Maric, 2008; Hyman, Ma, Balaguer-Ballester, Durstewitz, & Seamans, 2012; Jung, Qin, McNaughton, & Barnes, 1991; Powell & Redish, 2014; Ragozzino, Kim, Hassert, Minniti, & Kiang, 2003; Wallis, Anderson, & Miller, 2001; Young & Shapiro, 2009; Ma et al., 2016; Mante et al., 2013). Neural activity in dmPFC reflects changes in behavioral strategy (Rich & Shapiro, 2009; Durstewitz, Vittoz, Floresco, & Seamans, 2010; Karlsson, Tervo, & Karpova, 2012; Powell & Redish, 2016; De Falco et al., 2019), and dmPFC contributes to decision-making and generating goal-directed actions (Seamans, Floresco, & Phillips, 1995; Killcross & Coutureau, 2003; Matsumoto, Suzuki, & Tanaka, 2003; Matsumoto & Tanaka, 2004; Hok, Save, Lenck-Santini, & Poucet, 2005; St. Onge and Floresco, 2009), potentially by controlling internal simulations of possible actions and their outcomes (Hassabis & Maguire, 2009; van der Meer, Kurth-Nelson, & Redish, 2012; Wang, Cohen, & Voss, 2015).

On the other hand, while hippocampus (HPC) is traditionally thought to represent spatial location and to play a central role in spatial navigation (O'Keefe & Dostrovsky, 1971; O'Keefe & Nadel, 1978; Redish, 1999), the activity of neurons in HPC also reflect cognitive, non-spatial information. For example, hippocampal place fields in rats depend on where the rat intends to go (Wood, Dudchenko, Robitsek, & Eichenbaum, 2000; Ferbinteanu & Shapiro, 2003; Smith & Mizumori, 2006), are modulated by context (Ferbinteanu, Shirvalkar, & Shapiro, 2011; Griffin, Eichenbaum, & Hasselmo, 2007; Hasselmo & Eichenbaum, 2005; Kennedy & Shapiro, 2009; Zilli & Hasselmo, 2008), and sometimes even completely remap or modulate their firing rates

* Corresponding author at: 6-145 Jackson Hall, 321 Church St SE, Minneapolis, MN 55455, USA.
E-mail address: redish@umn.edu (A.D. Redish).

https://doi.org/10.1016/j.nlm.2020.107215
Received 6 January 2020; Received in revised form 7 March 2020; Accepted 18 March 2020
Available online 08 April 2020
1074-7427/

depending on sensory cues (Sharp, Kubie, & Muller, 1990; Sharp, Blair, Etkin, & Tzanetos, 1995; Leutgeb et al., 2005; Kubie, Fenton, Novikov, Touretzky, & Muller, 2007; Bahar, Shirvalkar, & Shapiro, 2011).

HPC and dmPFC, along with other structures, are thought to form a processing loop where bottom-up information from HPC informs representations in dmPFC, and top-down signals from dmPFC influence memory retrieval in HPC based on context or strategy (Wang et al., 2015; Jai & Frank, 2015; Redish, 2016; Shin & Jadhav, 2016; Eichenbaum, 2017). There are both direct and indirect projections from HPC to dmPFC (Swanson, 1981; Ferino, Thierry, & Glowinski, 1987; Jay & Witter, 1991; Verwer, Meijer, Van Uum, & Witter, 1997; Delatour & Witter, 2002; Floresco & Grace, 2003; Hoover & Vertes, 2007), as well as bidirectional connections between dmPFC and HPC via the nucleus reuniens of the thalamus and other thalamic nuclei (Vertes, 2002; Vertes, 2004; McKenna & Vertes, 2004; Vertes, Hoover, Do Valle, Sherman, & Rodriguez, 2006; Di Prisco & Vertes, 2006; Vertes, Hoover, Szigeti-Buck, & Leranth, 2007; Hoover & Vertes, 2012; Cassel et al., 2013; Dolleman-Van der Weel, Da Silva, & Witter, 2017; Dolleman-van der Weel et al., 2019). Contralateral lesion studies suggest that a PFC-HPC connection is required for spatial working memory (Floresco, Seamans, & Phillips, 1997; Wang & Cai, 2006; Wang & Cai, 2008; Barker et al., 2017; Maharjan, Dai, Glantz, & Jadhav, 2018), and electrophysiology experiments find interactions between HPC and dmPFC local field potentials, spike times, or both (Benchenane et al., 2010; Brincat & Miller, 2015; Colgin, 2011; Fujisawa & Buzsáki, 2011; Gordon, 2011; Hyman, Hasselmo, & Seamans, 2011; Hyman, Zilli, Paley, & Hasselmo, 2005, 2010; Jones & Wilson, 2005; Liu, Bai, Xia, & Tian, 2018; O'Neill, Gordon, & Sigurdsson, 2013; Place, Farovik, Brockmann, & Eichenbaum, 2016; Siapas, Lubenov, & Wilson, 2005; Sirota et al., 2008; Spellman et al., 2015; Zielinski, Shin, & Jadhav, 2019). Disrupting PFC changes information representation in HPC (Hok, Chah, Save, & Poucet, 2013; Navawongse & Eichenbaum, 2013; Guise & Shapiro, 2017; Schmidt, Duin, & Redish, 2019), an effect which is thought to be mediated by the nucleus reuniens (Dolleman-van der Weel, Morris, & Witter, 2009; Griffin, 2015; Hallock, Wang, & Griffin, 2016; Hallock, Wang, Shaw, & Griffin, 2013; Ito, Moser, & Moser, 2018; Layfield, Patel, Hallock, & Griffin, 2015; Linley, Gallo, & Vertes, 2016; Maisson, Gemzik, & Griffin, 2018; Mei, Logothetis, & Eschenko, 2018; Viena, Linley, & Vertes, 2018; Zimmerman & Grace, 2018; Ito et al., 2015; Xu and Südhof, 2013). Conversely, disrupting HPC outputs to the PFC weakens the encoding of spatial working memories (Spellman et al., 2015).

Do representations in dmPFC reflect new contingencies before, after, or at the same time as representations in HPC? Current theories which posit that dmPFC exerts a top-down contextual influence on HPC certainly lead to the hypothesis that updates in contingency or strategy representations would be seen in dmPFC before HPC. But even if representations of new rules in dmPFC do precede those in HPC, the time scale of this lead is difficult to deduce: it could be anywhere from tens of milliseconds to minutes. Guise and Shapiro (2017) examined the interaction between contingency representations by HPC and mPFC ensembles on a task with rule switches. They provided convincing evidence that both HPC and mPFC represent goals, and that some HPC activity is predictable from past mPFC activity on trials after rule changes. However, the study did not investigate the time course or latency of this interaction.

In addition to representing contingency or strategy information, ensemble activity in both HPC and PFC has been found to change slowly over time (Mankin et al., 2012; Hyman, Ma, Balaguer-Ballester, Durstewitz, & Seamans, 2012; Ziv et al., 2013; Malagon-Vina, Ciocchi, Passecker, Dorffner, & Klausberger, 2018; Mau et al., 2018). Some theories suggest that this change over time provides a form of temporal context (Howard & Kahana, 2002; Mensink and Raaijmakers, 1988), because events which occur close in time would have similar representations or "timestamps" (Rubin, Geva, Sheintuch, & Ziv, 2015), leading to easier retrieval of temporally similar memories. The change of ensemble encoding over time could also be due to the representation of other unmeasured factors such as motivation or satiation. While this representational change appears to occur in both HPC and dmPFC, it is unknown how quickly representations change in each structure relative to the other. Furthermore, drifting neural activity over time could be misconstrued as task rule representation when the rules are presented in blocks of time, and so it remains unclear how much of contingency or strategy representations in block-structured tasks can be explained simply by representational drift.

In this study, we developed an analysis to disambiguate the contributions of contingency representation from the representation of other non-contingency time-varying information. We determined that representations in dmPFC and HPC encode task contingencies while simultaneously changing over time in ways unrelated to contingency, and that contingency representation could not be explained by this encoding drift. We also compared the time course of contingency encoding changes between dmPFC and HPC as task rules changed, and compared the rate of change of representations in both structures.

We recorded neural ensembles simultaneously in rodent dmPFC and CA1 of dorsal hippocampus as rats performed a new variant of the multiple-T Left/Right/Alternate (MT-LRA) task (Gupta, van der Meer, Touretzky, & Redish, 2010; Blumenthal, Steiner, Seeland, & Redish, 2011; Steiner & Redish, 2012; Gupta, Van Der Meer, Touretzky, & Redish, 2012; Powell & Redish, 2014; Regier, Amemiya, & Redish, 2015; Powell & Redish, 2016), a spatial reversal task where rats were required to adjust their behavioral strategies after uncued contingency changes. We modified the task to include multiple uncued contingency changes per session (once per about 30 trials), whereas previous studies using this task employed only a single contingency switch per session. We found that both dmPFC and CA1 ensembles encoded the current contingency, and that this encoding in both dmPFC and CA1 transitioned to represent the new contingency after rule switches. However, we found that the transition in dmPFC preceded the transition in CA1: the transition in CA1 occurred at roughly the same time as animals updated their behavioral strategies to be consistent with the new contingency, while the transition in dmPFC preceded this behavioral change. Furthermore, because our modified task had multiple contingency blocks of the same type per session, we were able to use a representational stability analysis to determine that in addition to representing contingency, neural representations in dmPFC and CA1 simultaneously represented other information which changed slowly over time. The speed of this representational change was faster in dmPFC than in HPC. Taken together, our findings suggest that dmPFC and HPC work together to represent task contingencies, even while simultaneously encoding other information which changes over time, but that the integration of top-down contextual information from dmPFC into hippocampal representations is non-instantaneous.

2. Materials and methods

2.1. Subjects

Six FBNF-1 rats aged 8–14 months at the beginning of behavior were subjects for the experiment (4 male, 2 female), bred from Fischer and Brown Norway rats (Envigo). Rats were housed on a 12 h light–dark cycle, and experimental sessions were conducted at the same time each day during the light phase. All experimental and animal care procedures complied with US National Institutes of Health guidelines for animal care and were approved by the Institutional Animal Care and Use Committee and the University of Minnesota.

2.2. Training and behavior on the MT-LRA task

The Multiple-T Left/Right/Alternate (MT-LRA) task was a spatial reversal task where rats were required to adjust their behavioral strategies after uncued rule changes. The maze consisted of several low-cost

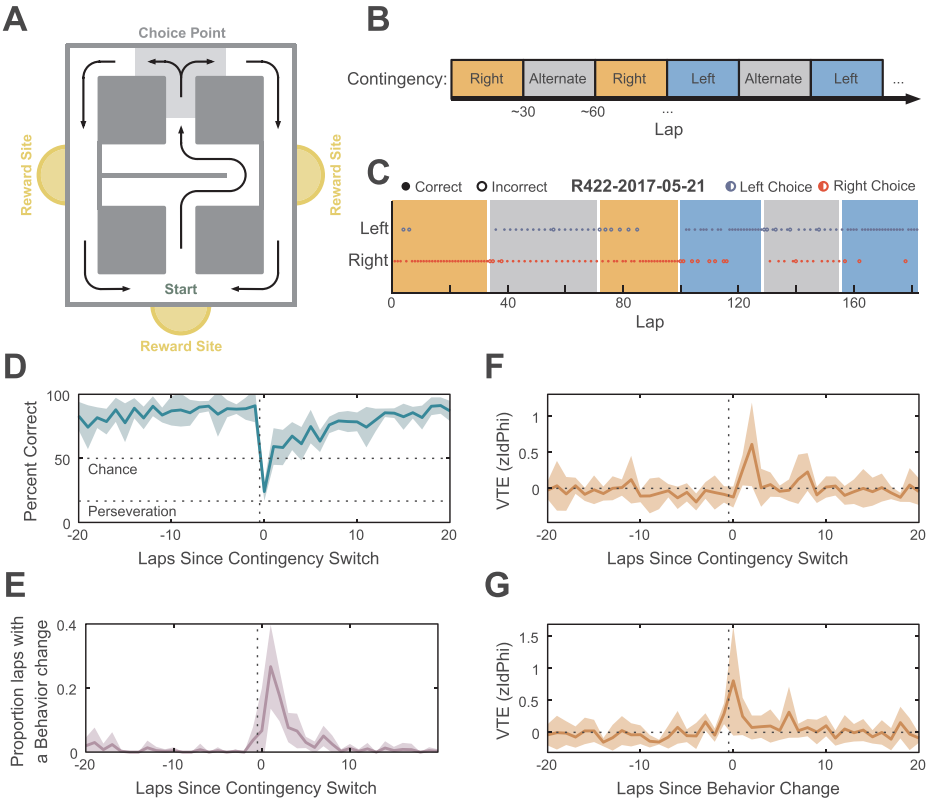

Fig. 1. Rat behavior on the MT-LRA task. (A) The MT-LRA task is a spatial maze with a choice point where rats receive rewards dependent on making choices consistent with the current contingency. (B) Contingencies are presented in blocks of laps lasting 30 ± 5 trials. (C) Example behavioral data from a single session. (D) Rat performance aligned to contingency switches. The vertical dotted line corresponds to the last lap of the previous contingency block. (E) Rat behavioral change laps aligned to contingency switches. The dotted line corresponds to the last lap of the previous contingency block. (F) VTE (measured by zIdPhi) aligned to contingency switches. (G) zIdPhi aligned to behavioral change laps. (D-G) show mean ± standard deviation, $N = 6$ rats.

choice points followed by a high-cost choice point between two actions: left or right (Fig. 1A). The walls of the maze were constructed using LEGO blocks, and the floor of the maze under the maze tracks was made up of white canvas. The configuration of the low-cost choice points at the center of the maze was determined by a single wall in the middle of the maze, which switched back and forth from the left to right side in a random manner at the start of a given day, after which it was held constant for the rest of the day. Each lap, if rats chose the action at the high-cost choice point which was consistent with the current contingency, they received one unflavored food pellet (45 mg each, Test Diet, Richmond, IN) at one of two reward sites on the side of the maze, and an additional food pellet at the rear of the maze. If their choice was inconsistent with the current contingency, no reward was delivered and rats had to circle around to the start of the maze to initiate a new lap. Two different auditory cues also signalled to the rats whether their decision was correct or incorrect: a swept-frequency sinewave "chirp" from 1 kHz to 3 kHz for correct, and two shorter 1 kHz square wave tones for incorrect. These auditory cues were presented as soon as rats entered the reward zone - at the same time as food delivery. After this moment, rats were unable to change their choice, and if they turned around to visit the opposite feeder, no food was delivered and rats had to initiate a new lap at the rear of the maze to continue the task. The contingency on any given lap was either Left (only left choices at the choice point led to reward), Right (only right choices), or Alternate (the opposite choice from the previous lap was required for reward).

The task used here differed from previous versions of the MT-LRA task in that the contingencies were presented in blocks: every 30 ± 5 laps, the contingency changed randomly to one of the other two contingencies. Rats were allowed to run laps freely on the task for one hour each day, and their daily food allowance came only from performing the task. However, rats were fed extra food after running the task if their weight dropped below 80% of their free-feeding weight.

Rat positions on the maze were tracked using a video camera (Cohu, Inc., San Diego, CA) placed above the maze. Custom Matlab (The MathWorks, Natick, MA) software determined animal position from the video, and controlled the state of the task (the current contingency, food pellet release, the presentation of audio cues, etc). The Matlab software also interfaced with an Arduino Uno Rev3 which ran custom software and triggered the release of food pellets from food pellet dispensers (MedAssociates, St. Albans, VT).

Rats were trained over the course of four weeks. Starting on the first week, rats were deprived of the freely available food in their home cages, but continued to have free access to water. Rats were handled and offered up to 15 g of food pellets each day for half an hour, to train them to eat the food pellets which would be available while performing the LRA task. For the second week, rats performed a simplified version of the task where the contingency was either Left or Right, and the contingency stayed constant throughout each session but changed randomly from session to session. Rats were rewarded with 2 food pellets per feeder at all feeder sites for the second week. For the third week, again there were no within-session contingency switches, but all three contingencies were possible (including Alternate), and only 1 food pellet was delivered at the rear food delivery site. For the final week of training, only 1 food pellet was delivered at all feeder sites, but the task was otherwise the same as during week 3.

After training, rats were given free access to food for at least 3 days, and then underwent surgery. After 3 days of post-surgery recovery, rats were again food deprived and re-trained for 1–2 weeks on the final training phase of the task (all 3 contingencies possible, but no within-session contingency switches, and 1 pellet per feeder). Finally, rats performed the full version of the task including within-session contingency switches and neural recordings for 2–3 weeks.

2.3. Surgery and electrode targets

After training, rats were given free access to food for at least 3 days, and then chronically implanted with a hyperdrive containing 24 tetrodes, and a separate drive containing a 32-site silicon probe (Cambridge NeuroTech, Cambridge, England). The hyperdrives contained two bundles of 12 tetrodes each, targeting the CA1 region of dorsal hippocampus bilaterally (3.8 mm posterior and ± 3.0 mm lateral from bregma). The hyperdrive for one rat contained a single bundle of 24 tetrodes targeting the right hippocampus. The silicon probes consisted of two 16-site shanks which were implanted 3.8 mm anterior to and 0.7 mm lateral from bregma at a 25 degree angle (targeting dmPFC on the right hemisphere, such that the final target was 2.3 mm A/P, −0.7 mm M/L, and 3.9 mm D/V, all coordinates relative to bregma). The hyperdrives and silicon probe drives were made in-house, and protective shrouds around the drives and amplifier boards were printed on a Form 2 3D printer (Formlabs, Somerville, MA).

Animals were anaesthetized with and maintained on isoflurane ($0.5 - 2\%$ isoflurane vaporized in O^2) for the duration of the surgery. Rats were placed in a sterotaxic apparatus (Kopf, Tujunga, CA) and were given penicillin (Combi-Pen-48) intramuscularly in each hindlimb, and carprofen (Rimadyl) subcutaneously. Rats' heads were shaved and disinfected with Betadine (Purdue Rederick, Norwalk, CT) before making an incision to reveal the skull. 3–5 jewlers' screws were used to anchor the drives to the skull, one of which was used as ground for the tetrodes, and a separate screw used as ground for the probe.

Three craniotomies were opened: two for the bilateral tetrode bundles using a surgical trephine, and one for the silicon probes using a burr. The dura was removed with forceps, the probe and tetrode drives were positioned with the sterotax, and then silicone gel (Dow Corning, Midland, MI) was applied to the craniotomies. A layer of MetaBond (Parkell, Edgewood, NY) and then dental acrylic (The Hygenic Corporation, Cuyahoga, OH) was applied to secure the drives to the skull. After surgery, the probes and tetrodes were turned down 640 μm. Rats were subcutaneously injected with carprofen on the day of surgery and for 2 days after surgery, as well as enrofloxacin (Enroflox) the day of surgery and for 5 days post-surgery.

2.4. Data acquisition and electrophysiology

Neural data from all rats was acquired on an Intan RHD2000 recording system (Intan Technologies, Los Angeles, CA), using four RHD2132 amplifier boards (three for the tetrodes and one for the silicon probe). The digitized signals were passed through a 24-channel commutator (Moog, East Aurora, NY) to allow the rats to move freely throughout recording sessions. To synchronize behavior with the neural recordings, timestamps were sent from the Matlab (Version 2017a, The MathWorks, Inc., Natick, MA) software running the task to digital input ports on the Intan RHD2000 USB Interface Board via an Arduino Uno.

Tetrodes were slowly advanced toward the hippocampal pyramidal layer, and the probes toward dmPFC, over the course of around 2 weeks, as the rats recovered and were re-trained on the task. The pyramidal layer was identified by the size of ripples and the direction of sharp wave deflection, as well as spike bursts during these SWRs.

Signals were filtered and spikes and LFP signals were extracted using in-house software written in Matlab and C. Spikes recorded on tetrodes in the hippocampus were manually clustered using the MClust 4.4 software package (Redish, 2017). Only well-separated clusters were kept and used for analysis. The median isolation distance was 21.2, and the median L-Ratio was 0.09 (Schmitzer-Torbert, Jackson, Henze, Harris, & Redish, 2005). Spikes recorded on silicon probes were sorted offline using Kilosort (Pachitariu, Steinmetz, Kadir, Carandini, & Harris, 2016) into putative clusters, and then manually refined using Phy (Rossant, Kadir, Goodman, Hunter, & Harris, 2016).

2.5. Histology

After rats were finished running the experiment, both tetrode and silicon probe recording locations were marked with electrolytic lesions. 10 μA was passed through a channel on each tetrode, and every fourth channel on the silicon probes, for 10s. At least two days after the lesions were made, the rats were anesthetized with a pentobarbital sodium solution (150 mg/kg, Fatal-Plus) and then perfused transcardially with saline followed by 10% formalin. Brains were stored in formalin, and then in a 30% sucrose formalin solution until slicing. Coronal slices were made through the hippocampus and prefrontal cortex (sagittal slices were made instead in PFC for 4 rats) using a cryostat, and the slices were stained with cresyl violet and imaged to determine tetrode and silicon probe recording locations (Fig. S5, aligned to (Paxinos and Watson, 2006)).

2.6. Statistics and general data analysis

All analyses were performed using Matlab (Version 2017a, The MathWorks, Inc., Natick, MA), except the binomial tests which were performed using the SciPy package for Python (Jones, Oliphant, & Peterson, 2019). For all neural analyses, we included only sessions where $\geqslant 10$ cells were recorded in dmPFC *and* $\geqslant 10$ cells were recorded in CA1 simultaneously.

2.7. Behavioral change point analysis

To identify laps where rats updated their behavioral choices to be consistent with the new contingency, we used a change point analysis from Gallistel, Fairhurst, and Balsam (2004). We considered 20 laps on either side of a contingency switch, after which the contingency in place was contingency X. We excluded laps which were before the previous switch, or after the next switch (in cases where contingency blocks lasted <20 laps). For each lap i in this window around each switch, we computed whether rats' choices were consistent with the new contingency (the rat made a choice which would be correct if X were the current contingency, $c_i = 1$) or inconsistent with the new contingency (the rat made a choice which would be incorrect if X were the current contingency, $c_i = 0$). We then applied the change point analysis from Gallistel et al. (2004) on $\mathbf{c}$ to determine on what lap rats were most likely to have updated their choices to be consistent with the new strategy.

2.8. Vicarious trial and error (VTE)

To measure vicarious trial and error behaviors, we used zIdPhi, a metric which captures both how long the rat hesitates at the choice point, and how quickly it moves its head back and forth (Papale, Stott, Powell, Regier, & Redish, 2012). When x and y are the x and y position of the rat's head,

$$\text{IdPhi} = \int_{\text{zone entry}}^{\text{zone exit}} \left| \frac{\delta}{\delta t} \text{atan2}\left(\frac{\delta y}{\delta t}, \frac{\delta x}{\delta t} \right) \right| \delta t$$

and zIdPhi is the IdPhi, z-scored within session.

To distinguish VTE events from non-VTE events (Fig. S4B), we fit a half-Gaussian distribution to values less than the mode of the zIdPhi distribution. We then assumed that zIdPhi values under a full Gaussian distribution with the same mean and standard deviation as the fit half-Gaussian corresponded to non-VTE events, and passes through the choice point with greater zIdPhi values corresponded to VTE events (Fig. S3B).

2.9. Representational stability analysis

To determine the extent to which ensemble representations

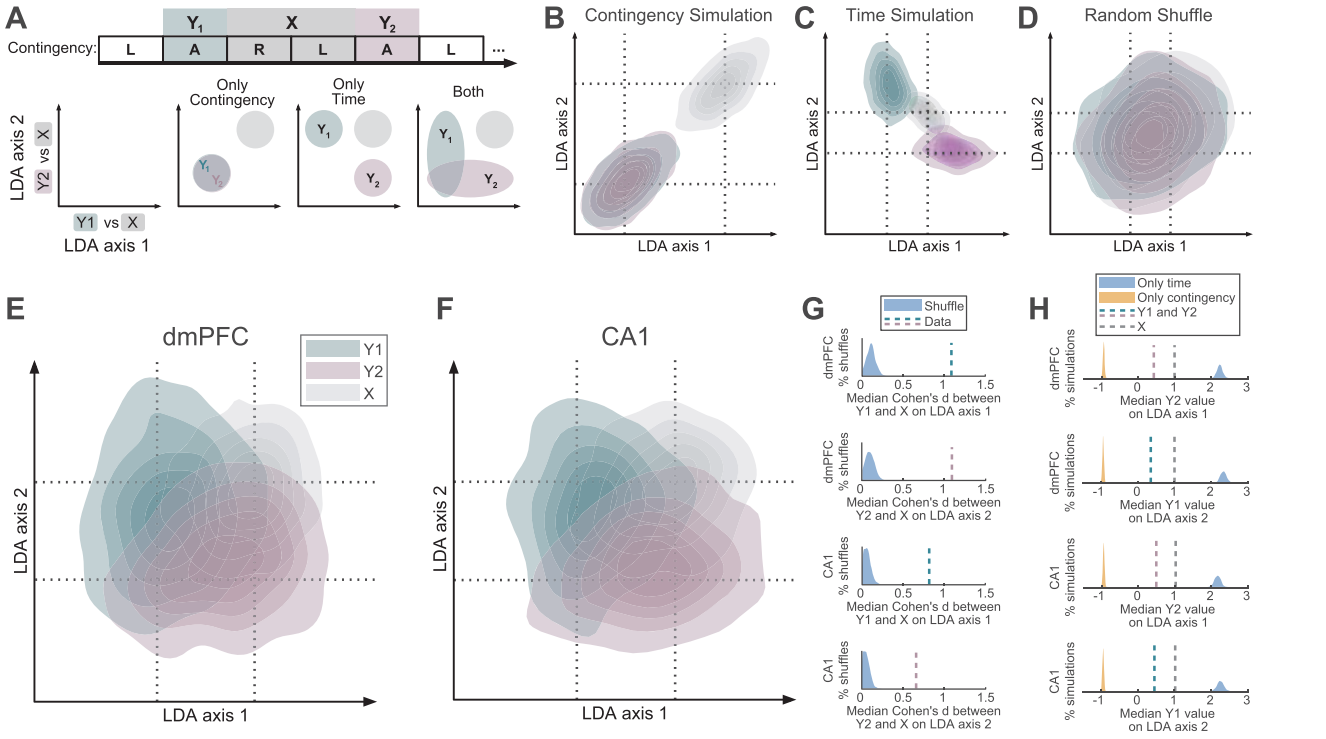

Fig. 2. Analysis of the stability of contingency representations. (A) Illustration of the representational stability analysis. (B-F) show kernel-density-smoothed contour density plots of the LDA projections of neural activity on each lap of each contingency pair used in the analysis (contingency pairs of the same type separated by block (s) of another type). $N = 62$ contingency epoch triplets, including a total of 4128 laps. Example simulations and shuffles in (B-D) used a matched number of contingency epoch triplets, laps, and units. (B) LDA projections of simulated firing rates containing only contingency information. (C) LDA projections of simulated firing rates containing only time information. (D) LDA projections of randomly shuffled neural activity in dmPFC. (E) LDA projections of neural activity in dmPFC, and (F) CA1. (G) Cohen's *d* between projections during *X* epochs and Y_1 epochs (green dotted bars) or Y_2 epochs (purple dotted bars), as compared to projections of shuffled neural activity (blue distributions). (H) Median LDA projection values for neural activity during Y_1 (green dotted bar), Y_2 (purple dotted bar), or *X* (gray dotted bar), as compared to simulations where firing rates contained only contingency information (orange distributions), or simulations where firing rates were controlled only by temporal drift (blue distributions). Dotted lines in (G) and (H) correspond to the median projection across the whole dataset, and each sample making up the distributions corresponds to the median of a shuffle or simulation with a matched number of cells, laps, and contingency triplets. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

reflected contingency encoding versus the encoding of temporal context (Fig. 2), we considered neural activity during three types of epochs: a contingency block of a given type (Y_1), a later contingency block of the same type (Y_2), and the time in between those two blocks (*X*). For each lap, we constructed a vector containing the firing rate of each neuron on that lap during the period during which the rat was within the central segment of the maze (from maze start to choice point entry). We used firing rates only during passes through the central segment to ensure that firing rate differences were not explainable by differences in location (if we had included firing rates over the entire maze, firing rates would likely have been trivially different between left and right laps, especially in HPC).

We used 20-fold cross-validated Linear Discriminant Analysis (LDA) to project these neural activity vectors onto the axes which best linearly separated samples during Y_1 from samples during *X* (the *x* axis in the projection plots in Fig. 2), and also onto the axes which separated samples during Y_2 from samples during *X* (the *y* axis in the projection plots in Fig. 2). Projections were normalized on both axes such that − 1 corresponded to the mean of projections during the *Y* epoch, and + 1 corresponded to the mean of the projections during the *X* epoch. We included only the last 20 trials of each contingency block, with the intent that this included laps only after rats had learned the true contingency.

2.10. Simulations and shuffles

To generate shuffled firing rates (Fig. 2D,G), we shuffled the inter-spike intervals for each cell independently, and re-generated that cell's spike times from the shuffled ISIs. This maintained some firing rate statistics (i.e. the mean firing rate of each cell) while removing any relationship to either contingency or time.

To generate firing rate simulations which represented only time and not contingency, we sorted the firing rates of each cell across a given session, such that a random half of the cells steadily increased their firing rates over the course of the session, and the other half decreased their firing rates over the course of the session. This ensured that we generated simulations which had the same number of cells and sessions as our actual data, as well as identical firing rate distributions as our actual data, but had firing rates which represented only the passage of time.

To generate firing rate simulations which represented only contingency and not time, we randomly assigned each cell a firing rate for each contingency, and added a small amount of noise. Again, this ensured that we generated "simulations" which had the same number of cells and sessions as our actual data, as well as identical firing rate distributions as our actual data, but had firing rates which represented only the current contingency.

2.11. Bayesian decoding of contingency

We used Bayesian decoding (Zhang, Ginzburg, McNaughton, & Sejnowski, 1998) to decode both spatial position and contingency from firing rates in either dmPFC or CA1 (Fig. 3). We used decoding time bins of 100 ms, and a 16 × 16 grid for spatial location, and 3 bins for the 3 contingencies. We used 100-fold cross validation to perform the decoding. The decoded contingency posterior probabilities were the

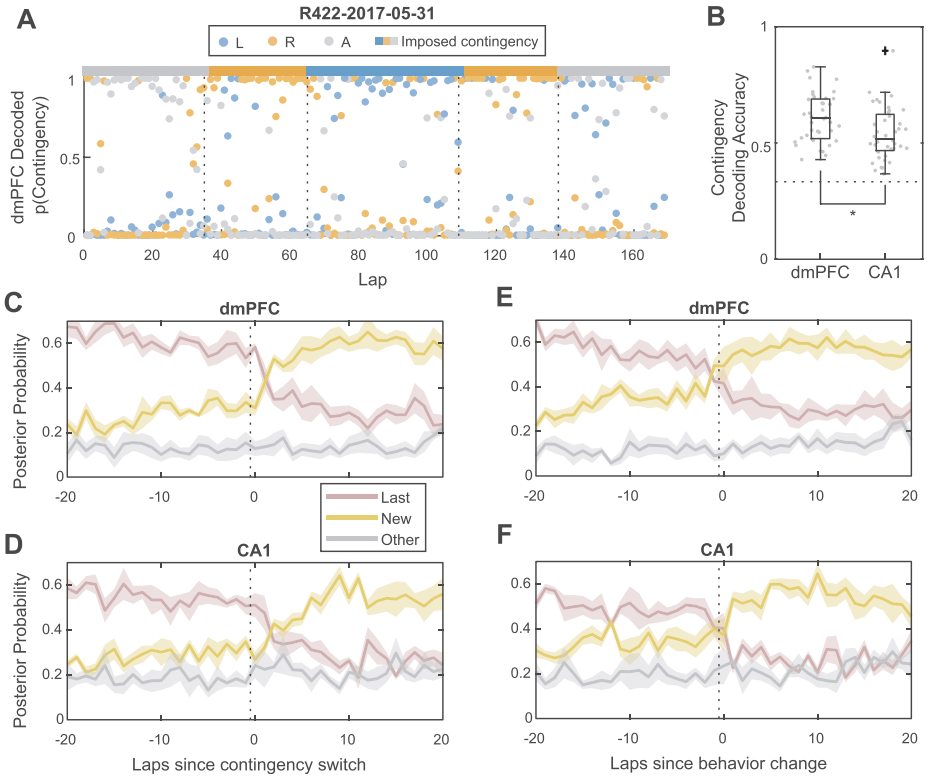

Fig. 3. Bayesian decoding of contingency representations in dmPFC and CA1. (A) Decoding from dmPFC ensemble activity over an example session. Dots are per-lap decoding posterior probabilities, and the colored bar at the top indicates the imposed contingency. (B) Decoding accuracy from ensembles in dmPFC and CA1. $N = 40$ sessions. Dotted line corresponds to chance. (C) Decoding aligned to contingency switches for dmPFC, and (D) CA1. Dotted lines indicate the last lap of the previous contingency block. (E) Decoding aligned to behavioral change laps for dmPFC, and (F) CA1. Dotted lines indicate the last lap of the previous behavioral strategy. (C-F) show mean ± SEM, $N = 6$ rats.

decoded posterior probabilities marginalized over spatial location. To compute the posterior probability of a given contingency on a given lap (as opposed to only during a single time bin), we computed the cumulative log posterior probability for each contingency across time samples during which the rat was on the central segment of the maze during that lap.

2.12. Ensemble firing rate correlations

To measure the similarity between population activity before and after contingency switches, we correlated neural tuning curves between pairs of laps around the switches (Fig. 4A,B). For each cell, we computed the firing rate in each of 5 spatial bins along the central segment of the maze (from lap start to just before choice point entry). For each pair of laps, we computed the Pearson correlation between the firing rate vectors for the lap (each of which had $5 \times N_c$ elements, where N_c is the number of cells recorded on that day). We averaged windows of 20 laps on either side of contingency switches to generate switch-aligned average correlation matrices for both dmPFC (Fig. 4A) and CA1 (Fig. 4B). We excluded laps before the previous switch or after the subsequent switch from this analysis, such that the data in Fig. 4A,B only reflected correlations between laps in identical or adjacent contingency blocks.

2.13. Ensemble transition point analysis

We performed a clustering-based transition point analysis to determine on what lap neural representations were most likely to transition from one representation to another (Powell & Redish, 2016). We used K-means clustering to cluster neural activity on each lap (again split into 5 spatial bins, only during the central segment of the maze, as in the correlation analysis above) into 3 clusters (for the 3 possible contingencies). We then applied the change point analysis from Gallistel et al. (2004) to the cluster IDs of each lap. Because K-means assigns clusters stochastically, we repeated this analysis 1000 times for each session to obtain a probability distribution across laps within the session, which captured the probability of neural activity transitioning from one state to another. These probability distributions were aligned to the switch to generate Fig. 4C.

To compute the timing difference between dmPFC and CA1 (Fig. 4D), we took the difference between the means of the transition probability distribution in dmPFC and CA1 for each switch, and computed bootstrapped 95% confidence intervals on the median.

2.14. Drift rate analysis

We determined the rate of representational drift by computing the rate at which the correlation between tuning curves decreased. We computed the mean firing rate within each contingency block for each cell in each of 5 spatial bins along the central segment of the maze, resulting in a vector with $5 \times N_c$ elements for each contingency block. Then, for each pair of contingency blocks within the same session which had identical types, we computed the Pearson correlation between their average firing rate vectors. To compute the speed of correlation decay, we divided the correlation coefficient by the number of laps which separated the centers of the two contingency blocks.

2.15. Ensemble remapping analysis

To compare the amount of global remapping occurring between contingency types, we measured the correlation of spatial tuning curves for the entire ensemble between contingency blocks. For each contingency block, we computed the average firing rate of each cell in each of 5 spatial bins along the central segment of the maze (from lap start,

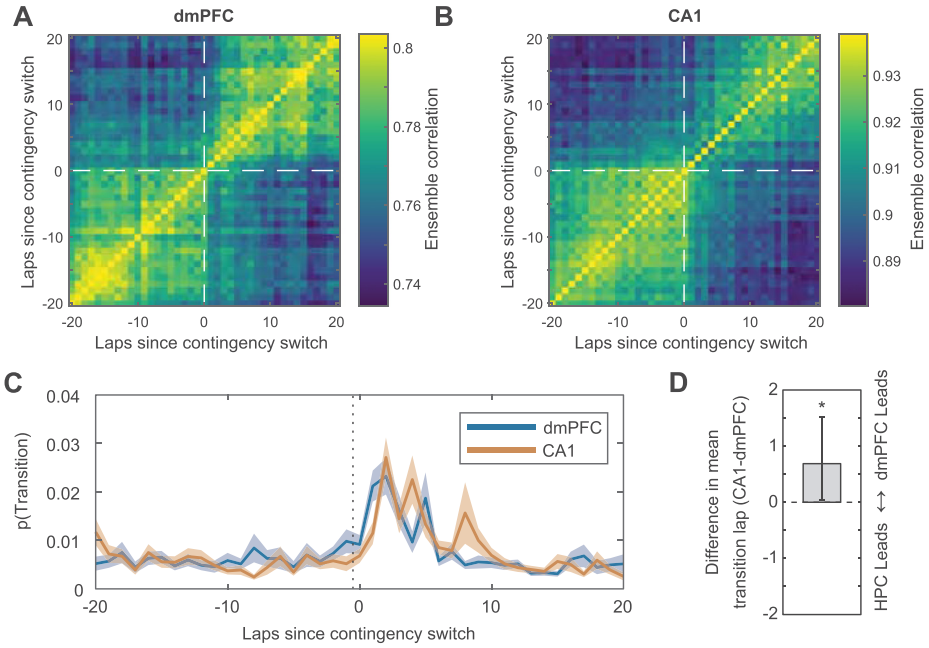

Fig. 4. Representational transitions after contingency switches in both dmPFC and CA1. (A) Switch-aligned correlation matrix of firing rates in dmPFC. (B) Switch-aligned correlation matrix of firing rates in CA1. (C) Transition probability for dmPFC and CA1 aligned to contingency switches. Lines and shaded areas show mean ± SEM, $N = 164$ switches. (D) Median difference of the transition probability distributions and bootstrapped 95% confidence interval. $N = 164$ switches.

at the rear of the maze, to choice point entry). We then took these per-contingency-block firing rate vectors and computed the Pearson's correlation coefficient between blocks (Fig. S10). Global remapping would have led to much lower correlation values between contingency blocks of a different type than between contingency blocks of the same type.

2.16. Cell stability analysis

To evaluate the stability of our electrophysiological recordings, we used an analysis similar to but not identical to that of Tolias et al. (2007). For each isolated unit we recorded, we split the spikes into those recorded in the first half of each session and those recorded during the second half, and took the average waveform for each epoch separately. We then computed the Euclidian distance of the average spike waveforms between the first and second halves of the session (blue distribution in Fig. S13). To generate a control distribution (the expected distribution of distances between waveforms which were not the same unit), we computed the distances between units and the average waveform of other units which were not identified as the same unit (the orange distribution in Fig. S13). Our identified units were substantially more self-similar than were non-identical units, suggesting our recordings did not suffer from excessive electrode drift.

3. Results

3.1. Rats updated behavioral strategies after contingency switches

We trained six rats to perform the MT-LRA task, in which rats ran through a spatial maze with one main choice point between two actions: left or right (Fig. 1A). On each lap, if rats chose the action at the choice point which was consistent with the current contingency, they received a food pellet reward. If their choice was inconsistent with the current contingency, no reward was delivered and rats had to circle around to the start of the maze to initiate a new lap. The contingency on any given lap was either Left (only left choices led to reward), Right (only right choices), or Alternate (the opposite choice from the previous lap was required for reward). The contingencies were presented in blocks: every 30 ± 5 laps, the contingency changed randomly to one of the other two contingencies (Fig. 1B). However, because of the contingency definitions, switches between all contingency types were not identical: switches from L or R blocks to any other type resulted in a 0% reward probability under perseveration, while switches from A to either L or R resulted in a 50% reward probability (in the case where the new contingency was consistent with the opposite of the choice the rat made on the previous lap, see Table S1).

Rats ran 137.7 ± 31.7 laps per session (mean ± standard deviation), and encountered 4.1 ± 1.2 contingency switches per session. Rats made correct choices (rewarded choices consistent with the current contingency) on 78.0 ± 3.0 percent of laps across all three contingency types, which was significantly more often than chance (4347 correct laps out of 5508, two sided binomial test vs 50%, $p < 10^{-100}$). Rats performed less well on laps during the Alternate contingency (Fig. S1C), where they made correct choices on only 62.9 ± 9.9 percent of laps, but their performance on laps during the Alternate contingency was still significantly better than chance (1231 correct laps out of 1874, two sided binomial test vs 50%, $p = 1.4 \times 10^{-42}$). Rats did not show any behavioral signs of anticipating the switch, as their choices did not reflect an increase in actions consistent with other contingencies as rats approached the expected contingency switch lap (Fig. S2).

The percentage of correct choices dropped on laps immediately following a contingency switch, but then increased over the course of the following contingency block, and plateaued well before the next contingency switch (Fig. 1D). A change-point analysis (Gallistel et al., 2004) indicated that rats updated their behavioral strategies to be consistent with the new contingency within about 5 laps of a contingency switch (Fig. 1E).

3.2. Rats deliberated over the course of the session, especially after contingency switches

While at choice points, rats sometimes display vicarious trial and error (VTE), a behavioral marker of deliberation (Redish, 2016). During VTE behaviors, rats pause and look back and forth down potential paths, as if deliberating over which path to choose (Fig. S3A). To quantify VTE, we measured *zIdPhi*, the z-scored integrated angular velocity of head movement (Papale et al., 2012). Note that VTE

behaviors are distinct from attentive scanning behaviors (Monaco, Rao, Roth, & Knierim, 2014), where rats rear and perform head-scanning movements, which are more exploratory in nature. As with many other studies which examine VTE (Schmidt, Papale, Redish, & Markus, 2013; Steiner & Redish, 2012; Stott & Redish, 2014), we observed low levels of zIdPhi on most choice point passes and higher levels of zIdPhi on fewer laps (Fig. S3B), suggesting rats deliberated on the minority of laps. A decrease in the amount of VTE over the course of a session is usually observed on other tasks (Papale et al., 2012; Breton, Seeland, & Redish, 2015; Redish, 2016), but on our task zIdPhi did not decrease over the course of the session (Fig. S3C). This suggests that the presence of multiple contingency switches, which continued to occur throughout the course of the session, repeatedly forced rats to deliberate and prevented them from fully automating their behavior on the task.

Although rats did not appear to automate their behavior over the course of an entire session, they did deliberate less as a given contingency block progressed. On laps immediately following a contingency switch, zIdPhi increased, and then decreased throughout the subsequent contingency block (Fig. 1F). This suggests that rats deliberated after contingency switches, but then automated as they learned the new contingency. However, this effect seemed to be mostly driven by switches to the Alternate contingency (Fig. S3D). The greatest levels of VTE were observed on laps where rats updated their behavioral strategies to be consistent with the new contingency (Fig. 1G, the median zIdPhi was significantly greater on laps where a behavioral change occurred than on other laps, two-sided Wilcoxon rank sum test, $p = 2.7 \times 10^{-5}$, $N = 164$ behavior change laps vs 5344 non-change laps).

Rats also displayed post-error slowing on our task (Laming, 1968; Narayanan & Laubach, 2008). Rats took significantly longer (around 1-2 s) to complete laps when the choice they made on the previous lap was incorrect (Fig. S4A, two-sided Wilcoxon rank sum test, $p = 0.031$, $N = 6$ rats). Post-error slowing was especially pronounced on laps where VTE occurred (Fig. S4B), which suggests that rats utilized a more conservative decision-making strategy following errors.

3.3. Ensemble activity represented contingency while simultaneously changing over time

We recorded simultaneously from neural ensembles in CA1 and dmPFC using two 16-site silicon probes (in dmPFC) and 24 tetrodes (in CA1) per animal. We limited our analyses to sessions where more than 10 cells were recorded in each structure simultaneously. During each of 40 behavioral sessions, we recorded 20.6 ± 5.4 units in dmPFC and 26.9 ± 13.5 units in CA1 (Fig. S6 and Table S2).

To determine the extent to which ensemble representations reflected contingency encoding versus the encoding of other information which changed over time, we took advantage of the fact that our task contained multiple contingency blocks of the same type within a session. We analyzed the stability of contingency representations across multiple presentations of the same contingency. Specifically, we examined ensemble activity during the first presentation of a contingency of a given type (contingency block Y_1, Fig. 2A), during a second presentation of that same contingency (contingency block Y_2), and during contingency blocks between the two (contingency block(s) X, during which the contingencies were of a different type). If time-varying information not related to contingency identity dominated the representations, ensemble activity during Y_1 and Y_2 should have been more dissimilar to each other than to ensemble activity during X, because they were further apart in time. On the other hand, if the contingency identities were encoded and other time-varying information did not dominate the representations, then ensemble activity during Y_1 and Y_2 should have been more similar to each other than to ensemble activity during X. Our dataset contained $N = 62$ of these contingency epoch triplets, including a total of 4128 laps.

We used cross-validated linear discriminant analysis (LDA) to project ensemble firing rates in dmPFC and CA1 onto the axis which best discriminated neural activity during Y_1 from that observed during X (Fig. 2A, on the horizontal axis), and also onto the axis which best discriminated neural activity during Y_2 from that observed during X (Fig. 2A, on the vertical axis).

To validate that this analysis was able to disambiguate contingency information from other time-varying information, we first applied the analysis to simulated data where firing rates represented only the contingency, and were otherwise time-invariant. Projections of these simulated firing rates during Y_1 and Y_2 overlapped, but were well-separated from projections of simulated firing rates during X (Fig. 2B). This is because the simulated activity during Y_1 and Y_2 were more similar to each other than they were to simulated activity during X (by design). We also applied this analysis to simulated data where firing rates represented only time, and did not explicitly represent the contingency. Projections of simulated firing rates representing only time information did not overlap at all, and projections during Y_1 and Y_2 were more dissimilar from each other than from projections during X (Fig. 2C). The projections of randomly shuffled neural activity had projections which extensively overlapped for all three epochs (Fig. 2D).

Projecting ensemble firing rates in dmPFC and CA1 onto LDA axes in this way revealed that neural activity in both structures showed signs of both contingency representation and a drift over time (Fig. 2E,F). The separation between projections during either Y_1 or Y_2 and X was significantly larger than would be expected by chance (Fig. 2G; 1000/1000 shuffles had lower Cohen's d between X and both Y projections, on both LDA axes in both structures). This suggests that neural encoding changed between subsequent contingency blocks. To determine whether this difference was due to an encoding of contingency identity, or simply due to an unrelated change in ensemble encoding over time, we compared the projections of neural activity to the simulations encoding only contingency information or only time information. LDA projections of neural activity during Y_1 and Y_2 were more similar to each other than the projections of simulations representing only time, and were more similar to each other than to the projections of neural activity during X epochs (Fig. 2H, 1000/1000 simulations had median projections further from the other Y block than was observed in the neural data for both LDA axes in both dmPFC and CA1). This suggests that the separation was in part due to representation of the contingency identity. However, not all of the representational change could be explained by contingency encoding: the LDA projections of neural activity during Y_1 and Y_2 were more separated than were the projections of simulations representing only contingency information (Fig. 2H, 1000/1000 simulations had median projections closer to the other Y block than was observed in the neural data for both LDA axes in both dmPFC and CA1). This indicates that while contingency identities were represented, there was also a change in the ensemble activity over time which could not be explained by contingency representation.

To determine if any single specific transition type was primarily responsible for these results, we repeated the LDA projection analysis for each of the six switch types individually, but found that the vast majority of projections for each individual switch type were similar to that of the pooled data (Fig. S7). Also, to account for whether the number of intervening contingency blocks was driving this effect, we repeated the LDA projection analysis for contingency block triplets containing only 1 intervening block, and those containing only 2 intervening blocks. The projections for each were again consistent with the projections when using all the data, though the projections for triplets with two intervening blocks were more separated, further suggesting the presence of a change over time (Fig.S8). These results demonstrate that ensemble activity in dmPFC and CA1 represented the abstract task rule or behavioral strategy, while simultaneously changing their representations over time in ways unrelated to contingency.

3.4. CA1 and dmPFC ensembles encoded the current contingency

To investigate how strongly dmPFC and HPC encoded the task rule, and when these representations changed, we used Bayesian decoding to decode the current task contingency (Left, Right, or Alternate) from ensemble firing rates in HPC or dmPFC. Unlike the previous analysis, this Bayesian decoding analysis captured both spatial information and contingency information, and did not simply depend on the average firing rate of cells across a trial. Both dmPFC and CA1 encoded the current contingency more strongly than chance (Fig. 3B, two-sided Wilcoxon signed rank test vs 1/3, $p = 3.6 \times 10^{-8}$ for both dmPFC and CA1, $N = 40$ sessions). 35 out of 40 sessions for dmPFC and 26 out of 40 sessions for CA1 had decoding accuracy greater than 3 standard deviations above the accuracy of decoding performed on shuffled firing rates (Fig. S9, 100 shuffles per session). Contingency decoding from ensembles in dmPFC was significantly more accurate than decoding from ensembles in CA1 (Fig. 3B, two-sided Wilcoxon rank sum test $p = 0.0045$, $N = 40$ sessions).

However, neither ensembles in CA1 nor dmPFC appeared to be overtly remapping between contingency types (Fig. S10), suggesting the encoding of contingency may have been via more subtle changes in firing rate, such as rate modulation (Leutgeb et al., 2005). Had ensembles remapped, correlations between ensemble firing rates between different contingency types would have been significantly less than the correlations between contingencies of the same type (Fig. S10).

Within around five laps after a contingency switch, ensembles in both dmPFC and CA1 began to represent the new contingency more strongly than the contingency from the previous block (Fig. 3C,D). On average, ensembles in dmPFC represented the new contingency more strongly than the previous contingency before the animal updated its behavior to be consistent with the new contingency (Fig. 3E), while in contrast this transition in CA1 was not different from the behavioral change point (Fig. 3F and S11B). Furthermore, ensemble firing rates in dmPFC and HPC were more correlated within contingency blocks than between contingency blocks, and the appearance of this within-contingency correlation appeared to occur more slowly after a contingency switch in CA1 than in dmPFC (Fig. 4A,B).

3.5. Ensembles in dmPFC transitioned to represent new contingencies before CA1 ensembles

To quantitatively measure the timing of these transitions on a switch-by-switch basis, we performed a change-point analysis on the ensemble firing rates in dmPFC and CA1. Briefly, we performed *k*-means clustering of normalized population firing rate vectors, and then applied a change point analysis (Gallistel et al., 2004) on the sequence of cluster IDs to determine on which laps neural activity transitioned from one cluster to another. The firing rate vector for each lap contained the average firing rate of each cell in each of 5 spatial bins across the central maze path from the beginning of a lap to the choice point. Because there is a level of stochasticity inherent in *k*-means clustering (a different initialization can result in a different clustering), we repeated the clustering and transition detection procedure many times to obtain a probability distribution of ensemble transitions over laps (Powell & Redish, 2016).

This analysis revealed that representations in both dmPFC and CA1 ensembles very likely underwent transitions within a few laps after contingency switches (Fig. 4C). The transition in dmPFC occurred significantly ahead of the transition in CA1 (paired two-sided Wilcoxon rank sum test comparing the means of the per-switch transition probability distributions, $p = 0.039$, $N = 160$ contingency switches). The median lead by dmPFC was 0.69 laps (Fig. 4D, 95% confidence interval between 0.045 and 1.5 laps). However note that our analysis was unable to resolve changes on sub-lap timescales, so we were only able to determine that the transition in dmPFC preceded that in CA1, but not by exactly how much.

In both dmPFC and CA1, the neural ensemble transition was most likely to occur on the lap when rats updated their behavioral strategies to be consistent with the new contingency, but dmPFC ensembles were significantly more likely to transition before the behavioral change, while the transition in CA1 did not occur at a significantly different time from the behavioral transition (Fig. S11B). There did not appear to be any major differences in the transition time course depending on the identity of the new contingency in either dmPFC or CA1 (Fig. S11C,D). However, in both dmPFC and CA1 there was a correlation between the amount of VTE and the probability of a neural transition (Fig. S12, two-sided Wilcoxon sign rank test $p = 0.020$ in dmPFC and $p = 0.0066$ in HPC, $N = 40$ sessions, of per session Spearman's rank correlation coefficients between zIdPhi and neural transition probabilities), suggesting that these brain areas were more likely to update their contingency representations on laps where rats deliberated.

3.6. Representations in both dmPFC and HPC also changed over time, independently of contingency

Comparing the projections of neural activity to simulations representing only time (Fig. 2) suggested that representations changed over time, but could not tell us how quickly. To investigate the speed at which neural representations in dmPFC and CA1 changed across time, without including changes due to the type of contingency, we measured the correlation of ensemble spatial tuning curves between pairs of contingency blocks of the same type, as a function of the time separating the blocks. Ensemble activity was less correlated between pairs of contingency blocks which were further apart in both dmPFC (Fig. 5A) and CA1 (Fig. 5B). The ensemble correlation coefficient decreased by 6.7×10^{-4} per lap in dmPFC (Fig. 5C, 95% confidence interval on the median was between -8.5×10^{-4} and -5.3×10^{-5}), while in CA1 the ensemble correlation coefficient decreased by 4.1×10^{-4} per lap (Fig. 5C, 95% confidence interval on the median was between -4.7×10^{-4} and -2.9×10^{-4}). The decrease per lap in the correlation coefficient was greater in dmPFC than in CA1 (Fig. 5C, Two-sided Wilcoxon rank sum test $p = 6.88 \times 10^{-5}$, $N = 62$ contingency pairs).

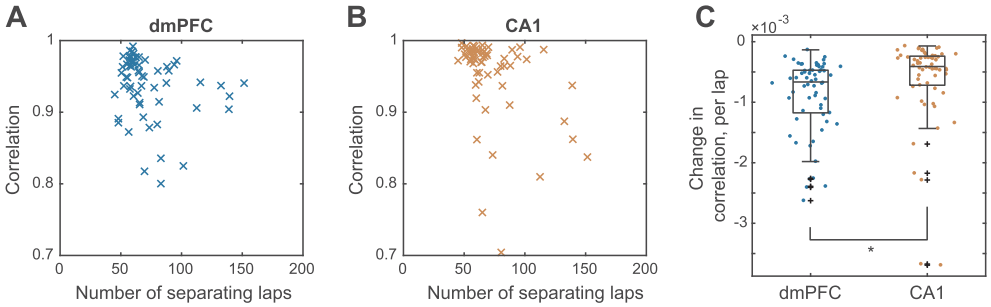

Fig. 5. The speed of the representational change over time. (A) Correlation between ensemble tuning curves as a function of how many laps separate the blocks being correlated, in dmPFC and (B) in CA1. (C) The change in the correlation per lap for both dmPFC and CA1. $N = 62$ pairs of contingency blocks of matching types.

This suggests that the change in the representation of non-contingency time-varying information, or representational drift, occurred more quickly in dmPFC than in CA1.

To validate that this change was not due simply to the physical drift of our recording electrodes over time, we compared the spike waveform self-similarity of identified units across time to the similarity between different units, similar to the analysis used by (Tolias et al., 2007) to validate stable recordings across days. We computed the Euclidean distance of spike waveforms between the first and second halves of the session during which that unit was recorded. We also computed the distances between the average waveforms of each unit and waveforms of other, non-identical, units within the same session. The distances between identified units across time was far lower than the distances between each unit and different units (Fig. S13). This indicates that our recordings suffered minimally from electrode drift, and thus electrode drift was unlikely to explain our observation of firing rate changes over time, suggesting that actual representational change accounted for the observed changes in firing rate patterns across time.

4. Discussion

Dorsomedial prefrontal cortex is thought to exert top-down contextual control on hippocampal spatial encoding. We recorded from ensembles in dmPFC and CA1 simultaneously on a task with multiple rule switches, and found that both dmPFC and CA1 represented task contingencies while concurrently representing other information which changed over time. Representations in dmPFC changed faster than in CA1, and dmPFC began to represent new task rules before HPC. Our results suggest that top-down information from dmPFC about contextual information, such as information about task rules, appears first in dmPFC and likely alters representations in HPC. The fact that the representational transition we observed in dmPFC preceded the transition in HPC suggests that this information may not immediately be incorporated into hippocampal representations, perhaps due to an inherent stability of hippocampal representations.

Given the fact that firing rates do change over time, when inspecting the neural encoding of information yoked to blocks of time, it becomes especially important to perform some sort of representational stability analysis (Fig. 2) to confirm that the decoding of this information is not simply an artifact of representational drift. The analysis we performed here is not the only option: any probabilistic classification model can be trained to separate neural activity during one presentation of a given contingency from activity between epochs, and then tested to see how well it classifies activity during a subsequent presentation of the contingency of interest. We opted for linear discriminant analysis both for its simplicity and because it lends itself to the more visually interpretable approach we took here (two dimensional projections, as opposed to comparing classification metrics).

Although representations are known to drift over time (Mankin et al., 2012; Hyman et al., 2012; Ziv et al., 2013), it is unknown what causes this drift, or what purpose it serves. One explanation is that the drift we observed could have been simply due to the representation of additional information which was changing over time, such as signals related to satiety or motivation. However, theories of how the brain encodes temporal context (Howard & Kahana, 2002; Howard et al., 2014; Liu et al., 2019; Mensink and Raaijmakers, 1988) suggest that drifting representations could facilitate the retrieval of memories which are closer in time, via associative dynamics. The temporal context model (Howard & Kahana, 2002) proposes that the drift is not random, but rather is driven by the retrieval of recent contextual information.

In a related study, Malagon-Vina et al. (2018) observed on a rule-switching task that strategy or rule representations in dmPFC differed between the first and later repeated rule presentations. They concluded that a new rule representation occurs each time a previously-presented rule occurs, suggesting that dmPFC may be encoding only rule changes, and not the actual rule identity. While it may appear that their conclusions are in conflict with our results that there are stable rule representations in dmPFC, we believe that our conclusions are actually consistent with those of Malagon-Vina et al. (2018). We observed that while there is a stable representation of rule or strategy identity, there is simultaneously a drift in the representation over time, due to any number of other factors (such as satiation or motivation). We hypothesize that this representational drift over time — and not any inherent change in contingency encoding — causes the apparent representation of rule identity to differ between the first and repeated presentations of the rule, which may have led Malagon-Vina et al. (2018) to conclude that a new representation is generated each time a rule is repeatedly presented. Of course this is somewhat of a semantic issue: the ensemble firing rates did indeed change between rule presentations. However, we find that even with the change in representation over time, there remains within dmPFC a consistent underlying representation of contingency.

While we did not test whether dmPFC and HPC were required for rats to perform our new variant of the task, we suspect that rats' performance on this task would be significantly disrupted after inhibition or lesioning dmPFC or HPC. Inactivation of or lesions in prelimbic cortex alone appears to impair rats' decisions when those decisions are context-dependent (Marquis et al., 2007; Dwyer et al., 2010). Previous work has also found that blocking NMDA receptors disrupts vicarious trial and error on the single-switch version of our task (Blumenthal et al., 2011), especially at the time of contingency switches. Since our new version of the task includes even more switches per session, we would expect then that disrupting the heavily NMDA-dependent PFC-hippocampal system would negatively impact rats' deliberation, leading to even poorer performance on the multi-switch version of the task. While previous studies on earlier versions of the contingency switch task have not found hippocampal map changes after contingency switches (Gupta et al., 2012), other studies have found changes in hippocampal representation related to context (Sharp et al., 1990; Hasselmo & Eichenbaum, 2005; Leutgeb et al., 2005; Kelemen & Fenton, 2010; Jezek, Henriksen, Treves, Moser, & Moser, 2011). More relevant to the experiments here, lesions of the fornix (which connects hippocampal and prefrontal areas, Jay & Witter (1991)) affect the ability to switch decision contexts (Hirsh, Leber, & Gillman, 1978).

While it is true that rats ran slower for a few laps following contingency switches, this post-error slowing is unable to explain the main findings in this paper. Both our Bayesian decoding of contingency (Fig. 3) and clustering-based transition analysis (Fig. 4) show observed changes in neural activity at the contingency change. If ensemble activity shifted as rats slowed, and then shifted back as they began to again run faster, we would expect to see two distinct peaks in the cluster transition likelihood of ensemble activity – a pattern not seen in our data (Fig. 4). Also, if a shift in speed was driving the change in neural activity, a Bayesian decoder would be unable to accurately decode contingency type. However, we found that a Bayesian decoder was able to decode the contingency type from neural activity in both dmPFC and CA1 (Fig. 3), even quickly after the contingency change. Most importantly, because we recorded ensemble activity simultaneously in dmPFC and CA1, a change in running speed which affected neural representations in these two areas would not explain why the transition in dmPFC occurred before that in CA1.

Our data shows that both dorsomedial prefrontal cortex and hippocampus encode contextual information about the current contingency, while simultaneously encoding other information which changes over time throughout the task. Furthermore, our results suggest that these context representations are more static in some brain areas, such as hippocampus, while they are more dynamic in others, such as dorsomedial prefrontal cortex.

Declaration of Competing Interest

The authors declare no competing interests.

Author Contributions

BH and ADR conceived and designed the experiment and analyses, and wrote the manuscript. BH collected and analyzed the data.

Acknowledgments

We thank Christopher Boldt, Kelsey Seeland, and Ayaka Sheehan for technical support and for building the tetrode drives, Christopher Boldt for building the silicon probe drives, Ayaka Sheehan for performing the histology, as well as Onni Rauhala and Daniel Min for help training rats. This work was funded by NSF IGERT Neuroengineering grant DGE-1069104, NIH R01-MH080318, and NIH R01-MH112688.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at https://doi.org/10.1016/j.nlm.2020.107215.

References

Bahar, A. S., Shirvalkar, P. R., & Shapiro, M. L. (2011). Memory-guided learning: Ca1 and ca3 neuronal ensembles differentially encode the commonalities and differences between situations. *Journal of Neuroscience, 31*(34), 12270–12281. https://doi.org/10.1523/JNEUROSCI.1671-11.2011.

Balleine, B. W., & Dickinson, A. (1998). Goal-directed instrumental action: Contingency and incentive learning and their cortical substrates. *Neuropharmacology, 37*(4–5), 407–419. https://doi.org/10.1016/S0028-3908(98)00033-1.

Barker, G. R. I., Banks, P. J., Scott, H., Ralph, G. S., Mitrophanous, K. A., Wong, L., Bashir, Z. I., Uney, J. B., & Warburton, E. C. (2017). Separate elements of episodic memory subserved by distinct hippocampal–prefrontal connections. *Nature Neuroscience, 20*(2), 242. https://doi.org/10.1038/nn.4472.

Benchenane, K., Peyrache, A., Khamassi, M., Tierney, P. L., Gioanni, Y., Battaglia, F. P., & Wiener, S. I. (2010). Coherent theta oscillations and reorganization of spike timing in the hippocampal-prefrontal network upon learning. *Neuron, 66*(6), 921–936. https://doi.org/10.1016/j.neuron.2010.05.013.

Blumenthal, A., Steiner, A., Seeland, K., & Redish, A. D. (2011). Effects of pharmacological manipulations of nmda-receptors on deliberation in the multiple-t task. *Neurobiology of Learning and Memory, 95*(3), 376–384. https://doi.org/10.1016/j.nlm.2011.01.011.

Breton, Y., Seeland, K. D., & Redish, A. D. (2015). Aging impairs deliberation and behavioral flexibility in inter-temporal choice. *Frontiers in Aging Neuroscience, 7*, 41. https://doi.org/10.3389/fnagi.2015.00041.

Brincat, S. L., & Miller, E. K. (2015). Frequency-specific hippocampal-prefrontal interactions during associative learning. *Nature Neuroscience, 18*(4), 576. https://doi.org/10.1038/nn.3954.

Cassel, J., De Vasconcelos, A. P., Loureiro, M., Cholvin, T., Dalrymple-Alford, J. C., & Vertes, R. P. (2013). The reuniens and rhomboid nuclei: Neuroanatomy, electrophysiological characteristics and behavioral implications. *Progress in Neurobiology, 111*, 34–52. https://doi.org/10.1016/j.pneurobio.2013.08.006.

Colgin, L. L. (2011). Oscillations and hippocampal–prefrontal synchrony. *Current Opinion in Neurobiology, 21*(3), 467–474. https://doi.org/10.1016/j.conb.2011.04.006.

Cowen, S. L., & McNaughton, B. L. (2007). Selective delay activity in the medial prefrontal cortex of the rat: Contribution of sensorimotor information and contingency. *Journal of Neurophysiology, 98*(1), 303–316. https://doi.org/10.1152/jn.00150.2007.

Dalley, J. W., Cardinal, R. N., & Robbins, T. W. (2004). Prefrontal executive and cognitive functions in rodents: Neural and neurochemical substrates. *Neuroscience & Biobehavioral Reviews, 28*(7), 771–784. https://doi.org/10.1016/j.neubiorev.2004.09.006.

De Falco, E., An, L., Sun, N., Roebuck, A. J., Greba, Q., Lapish, C. C., et al. (2019). The rat medial prefrontal cortex exhibits flexible neural activity states during the performance of an odor span task. *eNeuro, 6*(2), https://doi.org/10.1523/ENEURO.0424-18.2019.

Delatour, B., & Gisquest-Verrier, P. (1999). Lesions of the prelimbic–infralimbic cortices in rats do not disrupt response selection processes but induce delay-dependent deficits: Evidence for a role in working memory? *Behavioral Neuroscience, 113*(5), 941. https://doi.org/10.1037/0735-7044.113.5.941.

Delatour, B., & Witter, M. P. (2002). Projections from the parahippocampal region to the prefrontal cortex in the rat: Evidence of multiple pathways. *European Journal of Neuroscience, 15*(8), 1400–1407. https://doi.org/10.1046/j.1460-9568.2002.01973.x.

Di Prisco, G. V., & Vertes, R. P. (2006). Excitatory actions of the ventral midline thalamus (rhomboid/reuniens) on the medial prefrontal cortex in the rat. *Synapse, 60*(1), 45–55. https://doi.org/10.1002/syn.20271.

Dolleman-Van der Weel, M. J., Da Silva, F. H. L., & Witter, M. P. (2017). Interaction of nucleus reuniens and entorhinal cortex projections in hippocampal field ca1 of the rat. *Brain Structure and Function, 222*(5), 2421–2438. https://doi.org/10.1007/s00429-016-1350-6.

Dolleman-van der Weel, M. J., Griffin, A. L., Ito, H. T., Shapiro, M. L., Witter, M. P., Vertes, R. P., et al. (2019). The nucleus reuniens of the thalamus sits at the nexus of a hippocampus and medial prefrontal cortex circuit enabling memory and behavior. *Learning & Memory, 26*(7), 191–205. https://doi.org/10.1101/lm.048389.118.

Dolleman-van der Weel, M. J., Morris, R. G. M., & Witter, M. P. (2009). Neurotoxic lesions of the thalamic reuniens or mediodorsal nucleus in rats affect non-mnemonic aspects of watermaze learning. *Brain Structure and Function, 213*(3), 329–342. https://doi.org/10.1007/s00429-008-0200-6.

Durstewitz, D., Vittoz, N. M., Floresco, S. B., & Seamans, J. K. (2010). Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron, 66*(3), 438–448. https://doi.org/10.1016/j.neuron.2010.03.029.

Dwyer, D. M., Dunn, M. J., Rhodes, S. E. V., & Killcross, A. S. (2010). Lesions of the prelimbic prefrontal cortex prevent response conflict produced by action–outcome associations. *Quarterly Journal of Experimental Psychology, 63*(3), 417–424. https://doi.org/10.1080/17470210903411049.

Eichenbaum, H. (2017). Prefrontal–hippocampal interactions in episodic memory. *Nature Reviews Neuroscience, 18*(9), 547. https://doi.org/10.1038/nrn.2017.74.

Euston, D. R., Gruber, A. J., & McNaughton, B. L. (2012). The role of medial prefrontal cortex in memory and decision making. *Neuron, 76*(6), 1057–1070. https://doi.org/10.1016/j.neuron.2012.12.002.

Ferbinteanu, J., & Shapiro, M. L. (2003). Prospective and retrospective memory coding in the hippocampus. *Neuron, 40*(6), 1227–1239. https://doi.org/10.1016/S0896-6273(03)00752-9.

Ferbinteanu, J., Shirvalkar, P., & Shapiro, M. L. (2011). Memory modulates journey-dependent coding in the rat hippocampus. *Journal of Neuroscience, 31*(25), 9135–9146. https://doi.org/10.1523/JNEUROSCI.1241-11.2011.

Ferino, F., Thierry, A. M., & Glowinski, J. (1987). Anatomical and electrophysiological evidence for a direct projection from ammon's horn to the medial prefrontal cortex in the rat. *Experimental Brain Research, 65*(2), 421–426. https://doi.org/10.1007/bf00236315.

Floresco, S. B., Block, A. E., & Maric, T. L. (2008). Inactivation of the medial prefrontal cortex of the rat impairs strategy set-shifting, but not reversal learning, using a novel, automated procedure. *Behavioural Brain Research, 190*(1), 85–96. https://doi.org/10.1016/j.bbr.2008.02.008.

Floresco, S. B., & Grace, A. A. (2003). Gating of hippocampal-evoked activity in prefrontal cortical neurons by inputs from the mediodorsal thalamus and ventral tegmental area. *Journal of Neuroscience, 23*(9), 3930–3943. https://doi.org/10.1523/JNEUROSCI.23-09-03930.2003.

Floresco, S. B., Seamans, J. K., & Phillips, A. G. (1997). Selective roles for hippocampal, prefrontal cortical, and ventral striatal circuits in radial-arm maze tasks with or without a delay. *Journal of Neuroscience, 17*(5), 1880–1890. https://doi.org/10.1523/JNEUROSCI.17-05-01880.1997.

Fujisawa, S., & Buzsáki, G. (2011). A 4 hz oscillation adaptively synchronizes prefrontal, vta, and hippocampal activities. *Neuron, 72*(1), 153–165. https://doi.org/10.1016/j.neuron.2011.08.018.

Gallistel, C. R., Fairhurst, S., & Balsam, P. (2004). The learning curve: implications of a quantitative analysis. *Proceedings of the National Academy of Sciences, 101*(36), 13124–13131. https://doi.org/10.1073/pnas.0404965101.

Gordon, J. A. (2011). Oscillations and hippocampal–prefrontal synchrony. *Current Opinion in Neurobiology, 21*(3), 486–491. https://doi.org/10.1016/j.conb.2011.02.012.

Griffin, A. L. (2015). Role of the thalamic nucleus reuniens in mediating interactions between the hippocampus and medial prefrontal cortex during spatial working memory. *Frontiers in Systems Neuroscience, 9*, 29. https://doi.org/10.3389/fnsys.2015.00029.

Griffin, A. L., Eichenbaum, H., & Hasselmo, M. E. (2007). Spatial representations of hippocampal ca1 neurons are modulated by behavioral context in a hippocampus-dependent memory task. *Journal of Neuroscience, 27*(9), 2416–2423. https://doi.org/10.1523/JNEUROSCI.4083-06.2007.

Guise, K. G., & Shapiro, M. L. (2017). Medial prefrontal cortex reduces memory interference by modifying hippocampal encoding. *Neuron, 94*(1), 183–192. https://doi.org/10.1016/j.neuron.2017.03.011.

Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S., & Redish, A. D. (2010). Hippocampal replay is not a simple function of experience. *Neuron, 65*(5), 695–705. https://doi.org/10.1016/j.neuron.2010.01.034.

Gupta, A. S., Van Der Meer, M. A. A., Touretzky, D. S., & Redish, A. D. (2012). Segmentation of spatial experience by hippocampal theta sequences. *Nature Neuroscience, 15*(7), 1032. https://doi.org/10.1038/nn.3138.

Haddon, J. E., & Killcross, A. S. (2005). Medial prefrontal cortex lesions abolish contextual control of competing responses. *Journal of the Experimental Analysis of Behavior, 84*(3), 485–504. https://doi.org/10.1901/jeab.2005.81-04.

Haddon, J. E., & Killcross, S. (2006). Prefrontal cortex lesions disrupt the contextual control of response conflict. *Journal of Neuroscience, 26*(11), 2933–2940. https://doi.org/10.1523/JNEUROSCI.3243-05.2006.

Hallock, H. L., Wang, A., & Griffin, A. L. (2016). Ventral midline thalamus is critical for hippocampal–prefrontal synchrony and spatial working memory. *Journal of Neuroscience, 36*(32), 8372–8389. https://doi.org/10.1523/JNEUROSCI.0991-16.2016.

Hallock, H. L., Wang, A., Shaw, C. L., & Griffin, A. L. (2013). Transient inactivation of the thalamic nucleus reuniens and rhomboid nucleus produces deficits of a working-memory dependent tactile-visual conditional discrimination task. *Behavioral Neuroscience, 127*(6), 860. https://doi.org/10.1037/a0034653.

Hassabis, D., & Maguire, E. A. (2009). The construction system of the brain. *Philosophical Transactions of the Royal Society B: Biological Sciences, 364*(1521), 1263–1271. https://doi.org/10.1098/rstb.2008.0296.

Hasselmo, M. E., & Eichenbaum, H. (2005). Hippocampal mechanisms for the context-

dependent retrieval of episodes. *Neural Networks, 18*(9), 1172–1190. https://doi.org/10.1016/j.neunet.2005.08.007.
Hirsh, R., Leber, B., & Gillman, K. (1978). Fornix fibers and motivational states as controllers of behavior: A study stimulated by the contextual retrieval theory. *Behavioral Biology, 22*(4), 463–478. https://doi.org/10.1016/S0091-6773(78)92583-X.
Hok, V., Chah, E., Save, E., & Poucet, B. (2013). Prefrontal cortex focally modulates hippocampal place cell firing patterns. *Journal of Neuroscience, 33*(8), 3443–3451. https://doi.org/10.1523/JNEUROSCI.3427-12.2013.
Hok, V., Save, E., Lenck-Santini, P. P., & Poucet, B. (2005). Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. *Proceedings of the National Academy of Sciences, 102*(12), 4602–4607. https://doi.org/10.1073/pnas.0407332102.
Hoover, W. B., & Vertes, R. P. (2007). Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. *Brain Structure and Function, 212*(2), 149–179. https://doi.org/10.1007/s00429-007-0150-4.
Hoover, W. B., & Vertes, R. P. (2012). Collateral projections from nucleus reuniens of thalamus to hippocampus and medial prefrontal cortex in the rat: A single and double retrograde fluorescent labeling study. *Brain Structure and Function, 217*(2), 191–209. https://doi.org/10.1007/s00429-011-0345-6.
Horst, N. K., & Laubach, M. (2009). The role of rat dorsomedial prefrontal cortex in spatial working memory. *Neuroscience, 164*(2), 444–456. https://doi.org/10.1016/j.neuroscience.2009.08.004.
Howard, M. W., Kahana, M. J. (2002). A. distributed representation of temporal context. Journal of Mathematical Psychology, vol. 46, 3 (pp. 269–299). https://doi.org/10.1006/jmps.2001.1388.
Howard, M. W., MacDonald, C. J., Tiganj, Z., Shankar, K. H., Du, Q., Hasselmo, M. E., & Eichenbaum, H. (2014). A unified mathematical framework for coding time, space, and sequences in the hippocampal region. Journal of Neuroscience, vol. 34, 13 (pp. 4692–4707). https://doi.org/10.1523/JNEUROSCI.5808-12.2014.
Hyman, J. M., Hasselmo, M. E., & Seamans, J. K. (2011). What is the functional relevance of prefrontal cortex entrainment to hippocampal theta rhythms? *Frontiers in Neuroscience, 5*, 24. https://doi.org/10.3389/fnins.2011.00024.
Hyman, J. M., Ma, L., Balaguer-Ballester, E., Durstewitz, D., & Seamans, J. K. (2012). Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proceedings of the National Academy of Sciences, 109*(13), 5086–5091. https://doi.org/10.1073/pnas.1114415109.
Hyman, J. M., Zilli, E. A., Paley, A. M., & Hasselmo, M. E. (2005). Medial prefrontal cortex cells show dynamic modulation with the hippocampal theta rhythm dependent on behavior. *Hippocampus, 15*(6), 739–749. https://doi.org/10.1002/hipo.20106.
Hyman, J. M., Zilli, E. A., Paley, A. M., & Hasselmo, M. E. (2010). Working memory performance correlates with prefrontal-hippocampal theta interactions but not with prefrontal neuron firing rates. *Frontiers in Integrative Neuroscience, 4*, 2. https://doi.org/10.3389/neuro.07.002.2010.
Ito, H. T., Zhang, S., Witter, M. P., Moser, E. I., Moser, M. (2015). A prefrontal–thalamo–hippocampal circuit for goal-directed spatial navigation. *Nature*, vol. 522, 7554: 50. https://doi.org/10.1038/nature14396.
Ito, H. T., Moser, E. I., & Moser, M. (2018). Supramammillary nucleus modulates spike-time coordination in the prefrontal-thalamo-hippocampal circuit during navigation. *Neuron, 99*(3), 576–587. https://doi.org/10.1016/j.neuron.2018.07.021.
Jai, Y. Y., & Frank, L. M. (2015). Hippocampal–cortical interaction in decision making. *Neurobiology of Learning and Memory, 117*, 34–41. https://doi.org/10.1016/j.nlm.2014.02.002.
Jay, T. M., & Witter, M. P. (1991). Distribution of hippocampal ca1 and subicular efferents in the prefrontal cortex of the rat studied by means of anterograde transport of phaseolus vulgaris-leucoagglutinin. *Journal of Comparative Neurology, 313*(4), 574–586. https://doi.org/10.1002/cne.903130404.
Jezek, K., Henriksen, E. J., Treves, A., Moser, E. I., & Moser, M.-B. (2011). Theta-paced flickering between place-cell maps in the hippocampus. *Nature, 478*(7368), 246–249. https://doi.org/10.1038/nature10439.
Jones, M. W., & Wilson, M. A. (2005). Theta rhythms coordinate hippocampal–prefrontal interactions in a spatial memory task. *PLoS Biology, 3*(12), e402. https://doi.org/10.1371/journal.pbio.0030402.
Jung, M. W., Qin, Y., McNaughton, B. L., & Barnes, C. A. (1991). Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cerebral Cortex, 8*(5), 437–450. https://doi.org/10.1093/cercor/8.5.437.
Karlsson, M. P., Tervo, D. G. R., & Karpova, A. Y. (2012). Network resets in medial prefrontal cortex mark the onset of behavioral uncertainty. *Science, 338*(6103), 135–139. https://doi.org/10.1126/science.1226518.
Jones, E., Oliphant, T., & Peterson, P., et al. (2019). SciPy: Open source scientific tools for Python. Version 1.3.1, 2001–2019. http://www.scipy.org.
Kelemen, E., & Fenton, A. A. (2010). Dynamic grouping of hippocampal neural activity during cognitive control of two spatial frames. *PLoS Biology*, vol. 8, 6. https://doi.org/10.1371/journal.pbio.1000403.
Kennedy, P. J., & Shapiro, M. L. (2009). Motivational states activate distinct hippocampal representations to guide goal-directed behaviors. *Proceedings of the National Academy of Sciences, 106*(26), 10805–10810. https://doi.org/10.1073/pnas.0903259106.
Kesner, R. P., & Churchwell, J. C. (2011). An analysis of rat prefrontal cortex in mediating executive function. *Neurobiology of Learning and Memory, 96*(3), 417–431. https://doi.org/10.1016/j.nlm.2011.07.002.
Killcross, S., & Coutureau, E. (2003). Coordination of actions and habits in the medial prefrontal cortex of rats. *Cerebral Cortex, 13*(4), 400–408. https://doi.org/10.1093/cercor/13.4.400.
Kubie, J. L., Fenton, A., Novikov, N., Touretzky, D., & Muller, R. U. (2007). Changes in goal selection induced by cue conflicts are in register with predictions from changes in place cell field locations. *Behavioral Neuroscience, 121*(4), 751. https://doi.org/10.1037/0735-7044.121.4.751.
Laming, D. R. J. (1968). *Information theory of choice-reaction times.* Academic Press.
Layfield, D. M., Patel, M., Hallock, H., & Griffin, A. L. (2015). Inactivation of the nucleus reuniens/rhomboid causes a delay-dependent impairment of spatial working memory. *Neurobiology of Learning and Memory, 125*, 163–167. https://doi.org/10.1016/j.nlm.2015.09.007.
Leutgeb, S., Leutgeb, J. K., Barnes, C. A., Moser, E. I., McNaughton, B. L., & Moser, M. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science, 309*(5734), 619–623. https://doi.org/10.1126/science.1114037.
Linley, S. B., Gallo, M. M., & Vertes, R. P. (2016). Lesions of the ventral midline thalamus produce deficits in reversal learning and attention on an odor texture set shifting task. *Brain Research, 1649*, 110–122. https://doi.org/10.1016/j.brainres.2016.08.022.
Liu, T., Bai, W., Xia, M., & Tian, X. (2018). Directional hippocampal-prefrontal interactions during working memory. *Behavioural Brain Research, 338*, 1–8. https://doi.org/10.1016/j.bbr.2017.10.003.
Liu, Y., Tiganj, Z., Hasselmo, M. E., & Howard, M. W. A (2019). Neural microcircuit model for a scalable scale-invariant representation of time. *Hippocampus*, vol. 29, 3 (pp. 260–274). https://doi.org/10.1002/hipo.22994.
Ma, L., Hyman, J. M., Durstewitz, D., Phillips, A. G., & Seamans, J. K. (2016). A quantitative analysis of context-dependent remapping of medial frontal cortex neurons and ensembles. *Journal of Neuroscience*, vol. 36, 31 (pp. 8258–8272). https://doi.org/10.1523/JNEUROSCI.3176-15.2016.
Maharjan, D. M., Dai, Y. Y., Glantz, E. H., & Jadhav, S. P. (2018). Disruption of dorsal hippocampal–prefrontal interactions using chemogenetic inactivation impairs spatial learning. *Neurobiology of Learning and Memory, 155*, 351–360. https://doi.org/10.1016/j.nlm.2018.08.023.
Maisson, D. J. N., Gemzik, Z. M., & Griffin, A. L. (2018). Optogenetic suppression of the nucleus reuniens selectively impairs encoding during spatial working memory. *Neurobiology of Learning and Memory, 155*, 78–85. https://doi.org/10.1016/j.nlm.2018.06.010.
Malagon-Vina, H., Ciocchi, S., Passecker, J., Dorffner, G., & Klausberger, T. (2018). Fluid network dynamics in the prefrontal cortex during multiple strategy switching. *Nature Communications, 9*(1), 309. https://doi.org/10.1038/s41467-017-02764-x.
Mankin, E. A., Sparks, F. T., Slayyeh, B., Sutherland, R. J., Leutgeb, S., & Leutgeb, J. K. (2012). Neuronal code for extended time in the hippocampus. *Proceedings of the National Academy of Sciences, 109*(47), 19462–19467. https://doi.org/10.1073/pnas.1214107109.
Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, vol. 503, 7474 (pp. 78). https://doi.org/10.1038/nature12742.
Marquis, J., Killcross, S., & Haddon, J. E. (2007). Inactivation of the prelimbic, but not infralimbic, prefrontal cortex impairs the contextual control of response conflict in rats. *European Journal of Neuroscience, 25*(2), 559–566. https://doi.org/10.1111/j.1460-9568.2006.05295.x.
Matsumoto, K., Suzuki, W., & Tanaka, K. (2003). Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science, 301*(5630), 229–232. https://doi.org/10.1126/science.1084204.
Matsumoto, K., & Tanaka, K. (2004). The role of the medial prefrontal cortex in achieving goals. *Current Opinion in Neurobiology, 14*(2), 178–185. https://doi.org/10.1016/j.conb.2004.03.005.
Mau, W., Sullivan, D. W., Kinsky, N. R., Hasselmo, M. E., Howard, M. W., & Eichenbaum, H. (2018). The same hippocampal ca1 population simultaneously codes temporal information over multiple timescales. *Current Biology, 28*(10), 1499–1508. https://doi.org/10.1016/j.cub.2018.03.051.
McKenna, J. T., & Vertes, R. P. (2004). Afferent projections to nucleus reuniens of the thalamus. *Journal of Comparative Neurology, 480*(2), 115–142. https://doi.org/10.1002/cne.20342.
Mei, H., Logothetis, N. K., & Eschenko, O. (2018). The activity of thalamic nucleus reuniens is critical for memory retrieval, but not essential for the early phase of "off-line" consolidation. *Learning & Memory, 25*(3), 129–137. https://doi.org/10.1101/lm.047134.117.
Mensink, G. & Raaijmakers, J. G. (1988). A model for interference and forgetting. *Psychological Review*, vol. 95, 4 (pp. 434). https://doi.org/10.1037/0033-295X.95.4.434.
Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 24*(1), 167–202. https://doi.org/10.1146/annurev.neuro.24.1.167.
Monaco, J. D., Rao, G., Roth, E. D., & Knierim, J. J. (2014). Attentive scanning behavior drives one-trial potentiation of hippocampal place fields. *Nature Neuroscience, 17*(5), 725. https://doi.org/10.1038/nn.3687.
Narayanan, N. S., & Laubach, M. (2008). Neuronal correlates of post-error slowing in the rat dorsomedial prefrontal cortex. *Journal of Neurophysiology, 100*(1), 520–525. https://doi.org/10.1152/jn.00035.2008.
Navawongse, R., & Eichenbaum, H. (2013). Distinct pathways for rule-based retrieval and spatial mapping of memory representations in hippocampal neurons. *Journal of Neuroscience, 33*(3), 1002–1013. https://doi.org/10.1523/JNEUROSCI.3891-12.2013.
O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*. https://doi.org/10.1016/0006-8993(71)90358-1.
O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map.* Oxford University Press.
O'Neill, P., Gordon, J. A., & Sigurdsson, T. (2013). Theta oscillations in the medial prefrontal cortex are modulated by spatial working memory and synchronize with the hippocampus through its ventral subregion. *Journal of Neuroscience, 33*(35), 14211–14224. https://doi.org/10.1523/JNEUROSCI.2378-13.2013.
Pachitariu, M., Steinmetz, N. A., Kadir, S. N., Carandini, M., & Harris, K. D. (2016). Fast

and accurate spike sorting of high-channel count probes with kilosort. In *Advances in neural information processing systems* (pp. 4448–4456).
Papale, A. E., Stott, J. J., Powell, N. J., Regier, P. S., & Redish, A. D. (2012). Interactions between deliberation and delay-discounting in rats. *Cognitive, Affective, & Behavioral Neuroscience, 12*(3), 513–526. https://doi.org/10.3758/s13415-012-0097-7.
Paxinos, G., & Watson, C. (2006). The rat brain in stereotaxic coordinates: hard cover edition. Elsevier.
Place, R., Farovik, A., Brockmann, M., & Eichenbaum, H. (2016). Bidirectional prefrontal-hippocampal interactions support context-guided memory. *Nature Neuroscience, 19*(8), 992. https://doi.org/10.1038/nn.4327.
Powell, N. J., & Redish, A. D. (2014). Complex neural codes in rat prelimbic cortex are stable across days on a spatial decision task. *Frontiers in Behavioral Neuroscience, 8*, 120. https://doi.org/10.3389/fnbeh.2014.00120.
Powell, N. J., & Redish, A. D. (2016). Representational changes of latent strategies in rat medial prefrontal cortex precede changes in behaviour. *Nature Communications, 7*, 12830. https://doi.org/10.1038/ncomms12830.
Preston, A. R., & Eichenbaum, H. (2013). Interplay of hippocampus and prefrontal cortex in memory. *Current Biology, 23*(17), R764–R773. https://doi.org/10.1016/j.cub.2013.05.041.
Ragozzino, M. E., & Kesner, R. P. (1998). The effects of muscarinic cholinergic receptor blockade in the rat anterior cingulate and prelimbic/infralimbic cortices on spatial working memory. *Neurobiology of Learning and Memory, 69*(3), 241–257. https://doi.org/10.1006/nlme.1998.3823.
Ragozzino, M. E., Kim, J., Hassert, D., Minniti, N., & Kiang, C. (2003). The contribution of the rat prelimbic-infralimbic areas to different forms of task switching. *Behavioral Neuroscience, 117*(5), 1054. https://doi.org/10.1037/0735-7044.117.5.1054.
Redish, A. D. (1999). *Beyond the Cognitive Map: From Place Cells to Episodic Memory*. Cambridge MA: MIT Press.
Redish, A. D. (2016). Vicarious trial and error. *Nature Reviews Neuroscience, 17*(3), 147. https://doi.org/10.1038/nrn.2015.30.
Redish, A.D. Mclust. Version 4.4.07, 2017. <http://redishlab.neuroscience.umn.edu/MClust/MClust.html>.
Regier, P. S., Amemiya, S., & Redish, A. D. (2015). Hippocampus and subregions of the dorsal striatum respond differently to a behavioral strategy change on a spatial navigation task. *Journal of Neurophysiology, 114*(3), 1399–1416. https://doi.org/10.1152/jn.00189.2015.
Rich, E. L., & Shapiro, M. (2009). Rat prefrontal cortical neurons selectively code strategy switches. *Journal of Neuroscience, 29*(22), 7208–7219. https://doi.org/10.1523/JNEUROSCI.6068-08.2009.
Rossant, C., Kadir, S., Goodman, D., Hunter, M., & Harris, K. (2016). Phy. Version 1.0.9, 2016. <http://github.com/cortex-lab/phy>.
Rubin, A., Geva, N., Sheintuch, L., & Ziv, Y. (2015). Hippocampal ensemble dynamics timestamp events in long-term memory. *Elife, 4*, e12247. https://doi.org/10.7554/eLife.12247.
Schmidt, B., Duin, A. A., & Redish, A. D. (2019). Disrupting the medial prefrontal cortex alters hippocampal sequences during deliberative decision-making. *Journal of Neurophysiology*. https://doi.org/10.1152/jn.00793.2018.
St. Onge, J. R. & Floresco, S. B. (2009). Prefrontal cortical contribution to risk-based decision making. *Cerebral Cortex*, vol. 20, 8 (pp. 1816–1828). https://doi.org/10.1093/cercor/bhp250.
Schmidt, B., Papale, A., Redish, A. D., & Markus, E. J. (2013). Conflict between place and response navigation strategies: Effects on vicarious trial and error (vte) behaviors. *Learning & Memory, 20*(3), 130–138. https://doi.org/10.1101/lm.028753.112.
Schmitzer-Torbert, N., Jackson, J., Henze, D., Harris, K., & Redish, A. D. (2005). Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience, 131*(1), 1–11. https://doi.org/10.1016/j.neuroscience.2004.09.066.
Seamans, J. K., Floresco, S. B., & Phillips, A. G. (1995). Functional differences between the prelimbic and anterior cingulate regions of the rat prefrontal cortex. *Behavioral Neuroscience, 109*(6), 1063. https://doi.org/10.1037/0735-7044.109.6.1063.
Sharp, P. E., Blair, H. T., Etkin, D., & Tzanetos, D. B. (1995). Influences of vestibular and visual motion information on the spatial firing patterns of hippocampal place cells. *Journal of Neuroscience, 15*(1), 173–189. https://doi.org/10.1523/JNEUROSCI.15-01-00173.1995.
Sharp, P. E., Kubie, J. L., & Muller, R. U. (1990). Firing properties of hippocampal neurons in a visually symmetrical environment: contributions of multiple sensory cues and mnemonic processes. *Journal of Neuroscience, 10*(9), 3093–3105. https://doi.org/10.1523/JNEUROSCI.10-09-03093.1990.
Shin, J. D., & Jadhav, S. P. (2016). Multiple modes of hippocampal–prefrontal interactions in memory-guided behavior. *Current Opinion in Neurobiology, 40*, 161–169. https://doi.org/10.1016/j.conb.2016.07.015.
Siapas, A. G., Lubenov, E. V., & Wilson, M. A. (2005). Prefrontal phase locking to hippocampal theta oscillations. *Neuron, 46*(1), 141–151. https://doi.org/10.1016/j.neuron.2005.02.028.
Sirota, A., Montgomery, S., Fujisawa, S., Isomura, Y., Zugaro, M., & Buzsáki, G. (2008). Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. *Neuron, 60*(4), 683–697. https://doi.org/10.1016/j.neuron.2008.09.014.
Smith, D. M., & Mizumori, S. J. Y. (2006). Learning-related development of context-specific neuronal responses to places and events: The hippocampal role in context processing. *Journal of Neuroscience, 26*(12), 3154–3163. https://doi.org/10.1523/JNEUROSCI.3234-05.2006.
Spellman, T., Rigotti, M., Ahmari, S. E., Fusi, S., Gogos, J. A., & Gordon, J. A. (2015). Hippocampal–prefrontal input supports spatial encoding in working memory. *Nature, 522*(7556), 309. https://doi.org/10.1038/nature14445.
Steiner, A. P., & Redish, A. D. (2012). The road not taken: neural correlates of decision making in orbitofrontal cortex. *Frontiers in Neuroscience, 6*, 131. https://doi.org/10.3389/fnins.2012.00131.
Swanson, L. W. (1981). A direct projection from ammon's horn to prefrontal cortex in the rat. *Brain Research*, vol. 217, 1 (pp. 150–154). https://doi.org/10.1016/0006-8993(81)90192-X.
Stott, J. J., & Redish, A. D. (2014). A functional difference in information processing between orbitofrontal cortex and ventral striatum during decision-making behaviour. *Philosophical Transactions of the Royal Society B: Biological Sciences, 369*, 20130472. https://doi.org/10.1098/rstb.2013.0472.
Tolias, A. S., Ecker, A. S., Siapas, A. G., Hoenselaar, A., Keliris, G. A., & Logothetis, N. K. (2007). Recording chronically from the same neurons in awake, behaving primates. *Journal of Neurophysiology, 98*(6), 3780–3790. https://doi.org/10.1152/jn.00260.2007.
Tronel, S., & Sara, S. J. (2003). Blockade of nmda receptors in prelimbic cortex induces an enduring amnesia for odor–reward associative learning. *Journal of Neuroscience, 23*(13), 5472–5476. https://doi.org/10.1523/JNEUROSCI.23-13-05472.2003.
Urban, K. R., Layfield, D. M., & Griffin, A. L. (2014). Transient inactivation of the medial prefrontal cortex impairs performance on a working memory-dependent conditional discrimination task. *Behavioral Neuroscience, 128*(6), 639. https://doi.org/10.1037/bne0000020.
van der Meer, M., Kurth-Nelson, Z., & Redish, A. D. (2012). Information processing in decision-making systems. *The Neuroscientist, 18*(4), 342–359. https://doi.org/10.1177/1073858411435128.
Vertes, R. P. (2002). Analysis of projections from the medial prefrontal cortex to the thalamus in the rat, with emphasis on nucleus reuniens. *Journal of Comparative Neurology, 442*(2), 163–187. https://doi.org/10.1002/cne.10083.
Vertes, R. P. (2004). Differential projections of the infralimbic and prelimbic cortex in the rat. *Synapse, 51*(1), 32–58. https://doi.org/10.1002/syn.10279.
Vertes, R. P., Hoover, W. B., Do Valle, A. C., Sherman, A., & Rodriguez, J. J. (2006). Efferent projections of reuniens and rhomboid nuclei of the thalamus in the rat. *Journal of Comparative Neurology*, vol. 499, 5 (pp. 768–796). https://doi.org/10.1002/cne.21135.
Vertes, R. P., Hoover, W. B., Szigeti-Buck, K., & Leranth, C. (2007). Nucleus reuniens of the midline thalamus: Link between the medial prefrontal cortex and the hippocampus. *Brain Research Bulletin, 71*(6), 601–609. https://doi.org/10.1016/j.brainresbull.2006.12.002.
Verwer, R. W. H., Meijer, R. J., Van Uum, H. F. M., & Witter, M. P. (1997). Collateral projections from the rat hippocampal formation to the lateral and medial prefrontal cortex. *Hippocampus, 7*(4), 397–402. https://doi.org/10.1002/(SICI)1098-1063(1997)7:4<397::AID-HIPO5>3.0.CO;2-G.
Viena, T. D., Linley, S. B., & Vertes, R. P. (2018). Inactivation of nucleus reuniens impairs spatial working memory and behavioral flexibility in the rat. *Hippocampus, 28*(4), 297–311. https://doi.org/10.1002/hipo.22831.
Wallis, J. D., Anderson, K. C., & Miller, E. K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature, 411*(6840), 953. https://doi.org/10.1038/35082081.
Wang, G., & Cai, J. (2006). Disconnection of the hippocampal–prefrontal cortical circuits impairs spatial working memory performance in rats. *Behavioural Brain Research, 175*(2), 329–336. https://doi.org/10.1016/j.bbr.2006.09.002.
Wang, G., & Cai, J. (2008). Reversible disconnection of the hippocampal-prelimbic cortical circuit impairs spatial learning but not passive avoidance learning in rats. *Neurobiology of Learning and Memory, 90*(2), 365–373. https://doi.org/10.1016/j.nlm.2008.05.009.
Wang, J. X., Cohen, N. J., & Voss, J. L. (2015). Covert rapid action-memory simulation (crams): A hypothesis of hippocampal–prefrontal interactions for adaptive behavior. *Neurobiology of Learning and Memory, 117*, 22–33. https://doi.org/10.1016/j.nlm.2014.04.003.
Wood, E. R., Dudchenko, P. A., Robitsek, R. J., & Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron, 27*(3), 623–633. https://doi.org/10.1016/S0896-6273(00)00071-4.
Xu, W., Südhof, T.C. (2013). A neural circuit for memory specificity and generalization. *Science*, vol. 339, 6125 (pp. 1290–1295). https://doi.org/10.1126/science.1229534.
Yoon, T., Okada, J., Jung, M. W., & Kim, J. J. (2008). Prefrontal cortex and hippocampus subserve different components of working memory in rats. *Learning & Memory, 15*(3), 97–105. https://doi.org/10.1101/lm.850808.
Young, J. J., & Shapiro, M. L. (2009). Double dissociation and hierarchical organization of strategy switches and reversals in the rat pfc. *Behavioral Neuroscience, 123*(5), 1028. https://doi.org/10.1037/a0016822.
Zhang, K., Ginzburg, I., McNaughton, B. L., & Sejnowski, T. J. (1998). Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells. *Journal of Neurophysiology, 79*(2), 1017–1044. https://doi.org/10.1152/jn.1998.79.2.1017.
Zielinski, M. C., Shin, J. D., & Jadhav, S. P. (2019). Coherent coding of spatial position mediated by theta oscillations in the hippocampus and prefrontal cortex. *Journal of Neuroscience, 39*(23), 4550–4565. https://doi.org/10.1523/JNEUROSCI.0106-19.2019.
Zilli, E. A., & Hasselmo, M. E. (2008). Modeling the role of working memory and episodic memory in behavioral tasks. *Hippocampus, 18*(2), 193–209. https://doi.org/10.1002/hipo.20382.
Zimmerman, E. C., & Grace, A. A. (2018). Prefrontal cortex modulates firing pattern in the nucleus reuniens of the midline thalamus via distinct corticothalamic pathways. *European Journal of Neuroscience, 48*(10), 3255–3272. https://doi.org/10.1111/ejn.14111.
Ziv, Y., Burns, L. D., Cocker, E. D., Hamel, E. O., Ghosh, K. K., Kitch, L. J., El Gamal, A., & Schnitzer, M. J. (2013). Long-term dynamics of ca1 hippocampal place codes. *Nature Neuroscience, 16*(3), 264. https://doi.org/10.1038/nn.3329.

Dorsomedial prefrontal cortex and hippocampus represent strategic context even while simultaneously changing representation throughout a task session

Brendan M. Hasz
A. David Redish

April 12, 2020

Published in *Neurobiology of Learning and Memory*.

SUPPLEMENTAL MATERIAL

Contingency switch type	Reward rate under perseveration
L $\rightarrow$ R	0%
L $\rightarrow$ A	0%
R $\rightarrow$ L	0%
R $\rightarrow$ A	0%
A $\rightarrow$ L	50%
A $\rightarrow$ R	50%

Table S1: (Related to Figure 1) The reward probability under perseveration (making actions consistent with the old contingency type) for different contingency switch types.

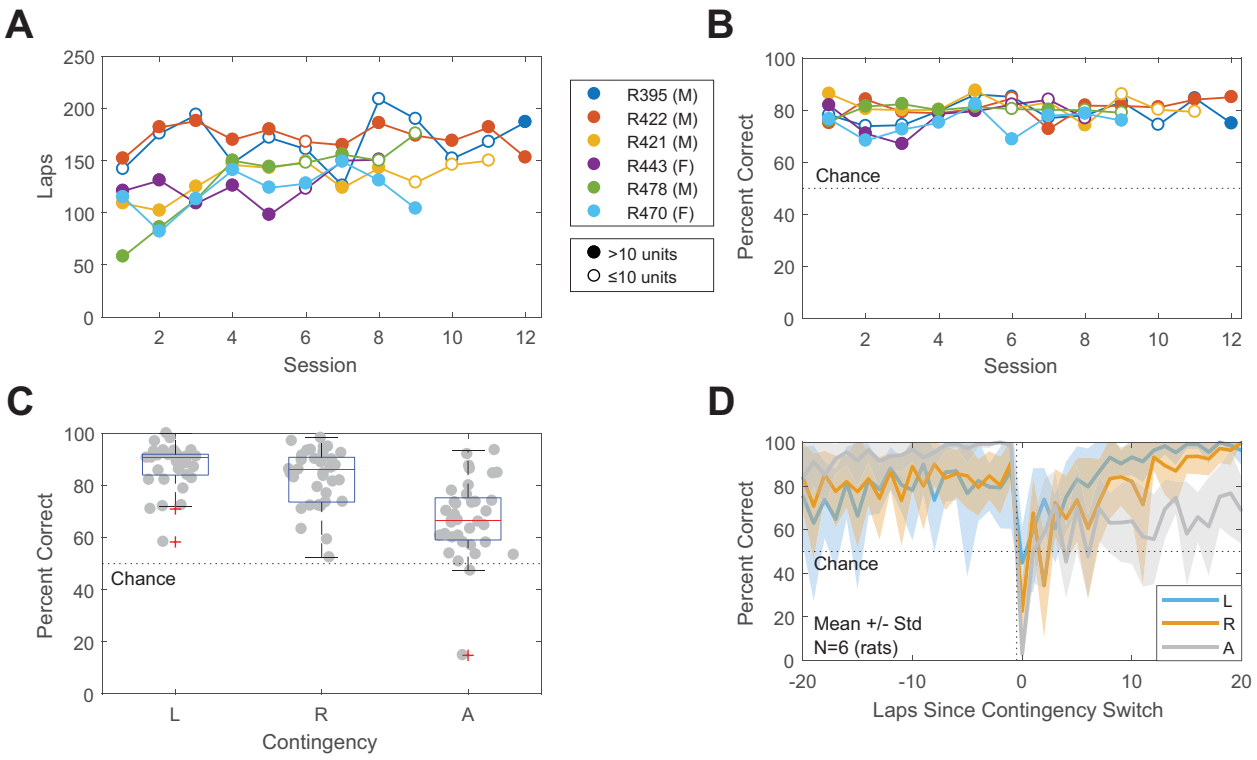

Figure S1: (Related to Figure 1) Behavioral performance on the MT-LRA task. (A) Laps per session ($N = 6$ rats). Filled circles indicate sessions which met the inclusion criteria (> 10 cells simultaneously recorded in both structures), and empty circles correspond to sessions which were not used for neural analyses. (B) Percent correct by session. (C) Performance by contingency ($N = 40$ sessions). (D) Performance aligned to switch, split by contingency.

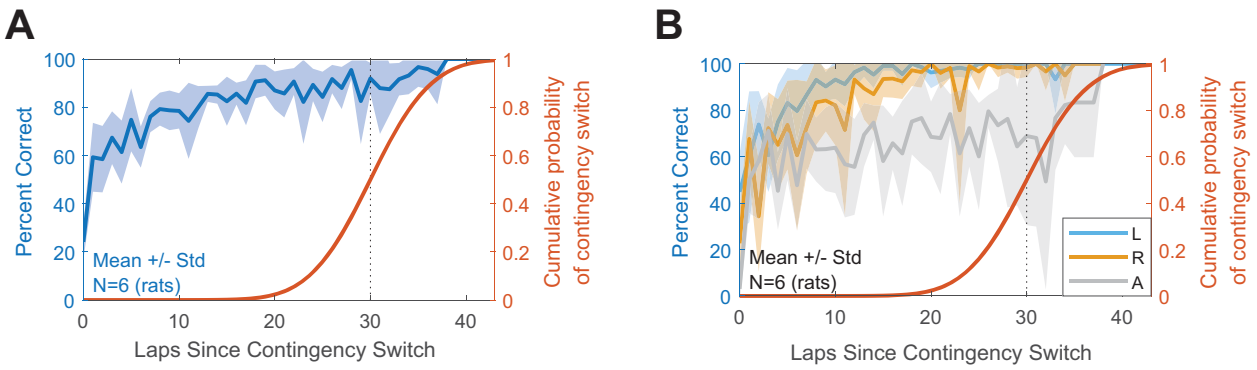

Figure S2: (Related to Figure 1) Performance aligned to the previous switch (A) overall, and (B) split by contingency.

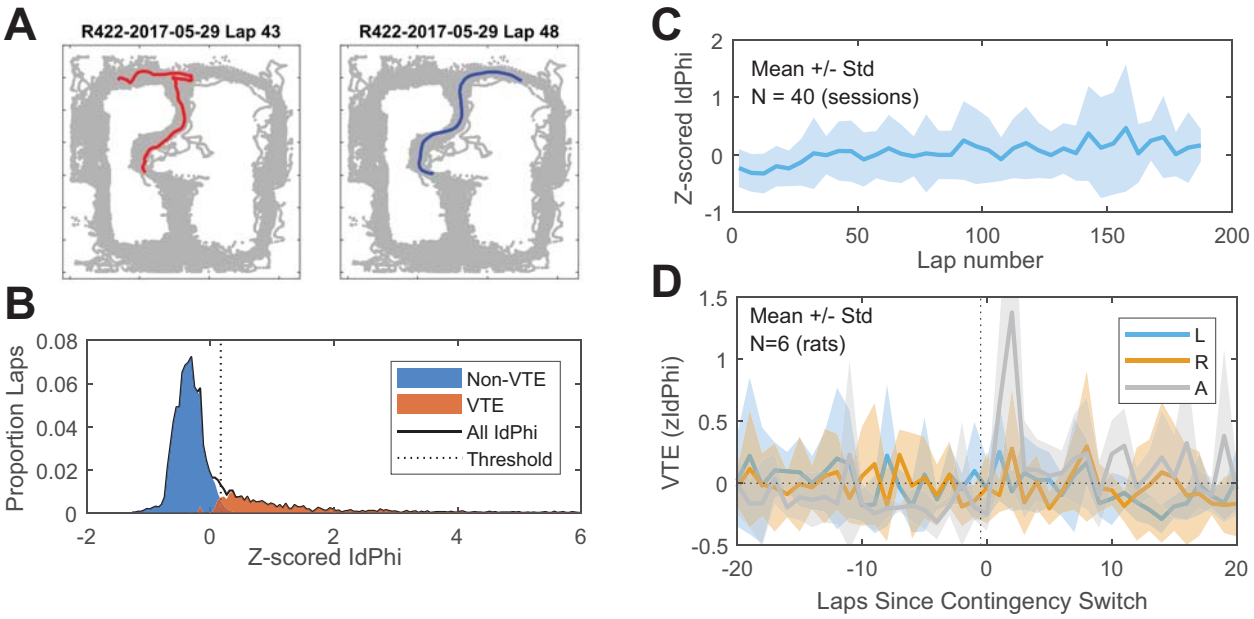

Figure S3: (Related to Figure 1) Vicarious trial and error (VTE) on the MT-LRA task. (A) Example of a pass through the choice point where the rat displayed VTE (left) and a non-VTE pass (right). (B) Distribution of zIdPhi across all laps and rats. (C) zIdPhi over the course of a session. (D) zIdPhi aligned to switch split by contingency.

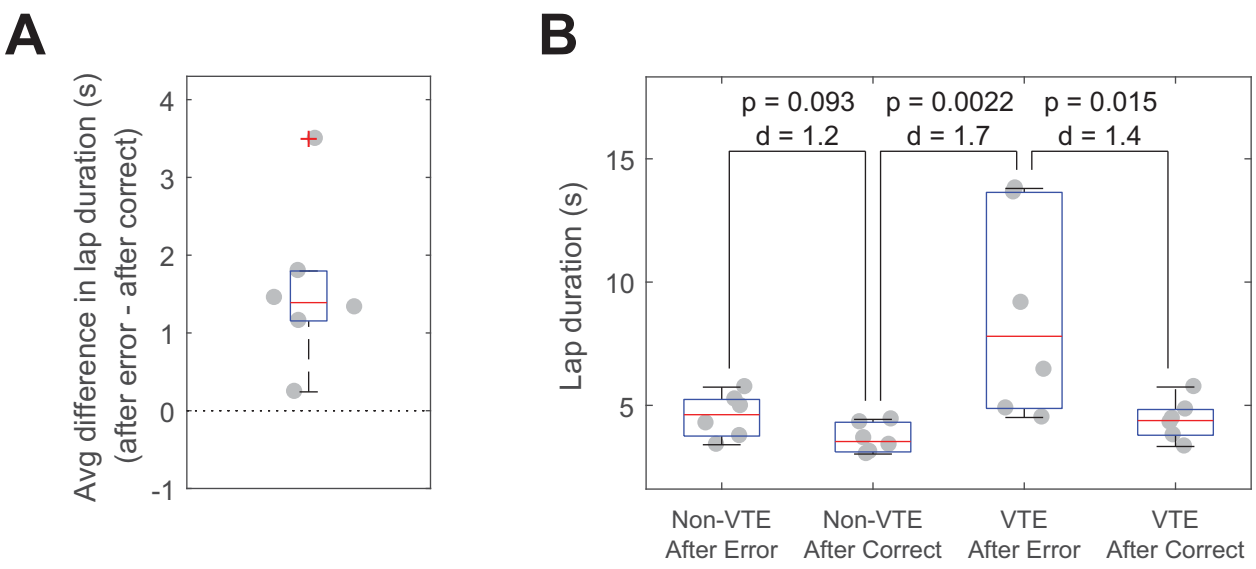

Figure S4: (Related to Figure 1) Post-error slowing. (A) The average difference in lap duration between laps following errors vs correct choices ($N = 6$ rats). (B) Lap duration split by both VTE and whether the rat made an error on the previous lap ($N = 6$ rats). Also shown are p values of two-sided Wilcoxon signed-rank tests, and Cohen's d.

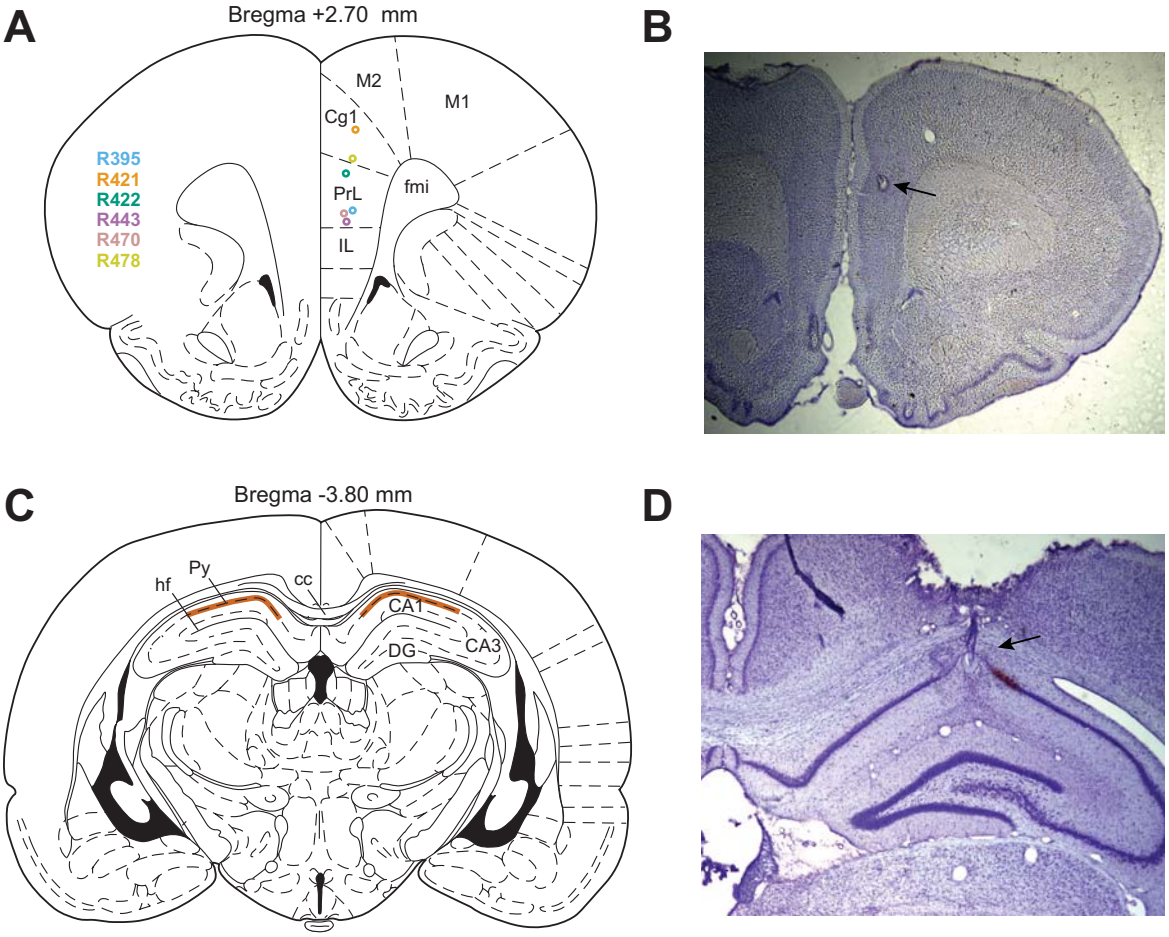

Figure S5: (Related to Figure 3) Histology and electrode targeting. (A) Si probe recording locations in dmPFC for each rat. (B) Example photo of a cresyl violet stained coronal slice through PFC, showing the electrolytic lesion created to mark the recording location (in Cg1 in this example). (C) Tetrode recording region (highlighted) was in the pyramidal layer of dorsal CA1. (D) Example photo of a cresyl violet stained coronal slice through HPC, showing an electrolytic lesion and electrode track. Anatomy diagrams in (A) and (C) are from Paxinos and Watson (2006). *IL*, infralimbic cortex. *PrL*, prelimbic cortex. *Cg1*, cingulate cortex area 1. *M2*, secondary motor cortex. *M1*, primary motor cortex. *fmi*, forceps minor of the corpus callosum. *hf*, hippocampal fissure. *Py*, pyramidal layer. *cc*, corpus callosum. *DG*, dentate gyrus.

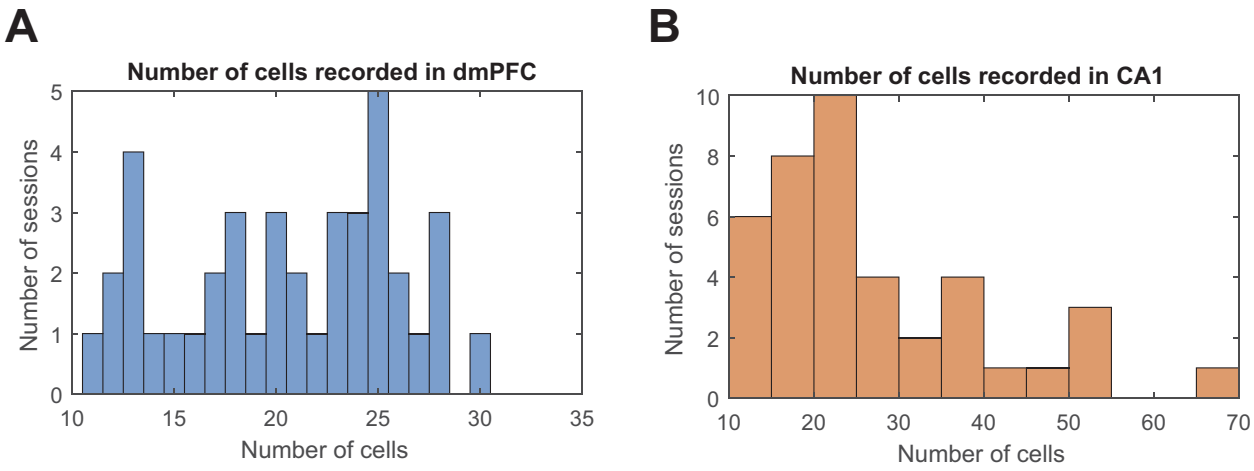

Figure S6: (Related to Figure 2) Number of cells recorded per session in (A) dmPFC and (B) in CA1.

Session	Rat	Sex	dmPFC Units	CA1 Units
1	1	M	12	13
2	2	M	12	39
3	2	M	23	19
4	2	M	22	11
5	2	M	20	20
6	2	M	20	24
10	2	M	18	36
11	2	M	18	27
12	2	M	21	11
13	3	M	28	38
14	3	M	25	19
15	3	M	28	18
16	3	M	25	15
17	3	M	15	23
18	3	M	13	19
19	3	M	14	19
20	3	M	13	20
21	4	F	16	29
22	4	F	13	24
23	4	F	17	24
24	4	F	11	23
25	4	F	13	22
26	5	M	24	31
27	5	M	23	30
28	5	M	24	38
29	5	M	24	21
30	5	M	28	11
31	5	M	18	14
32	6	F	25	22
33	6	F	26	48
34	6	F	21	53
35	6	F	19	51
36	6	F	25	70
37	6	F	26	51
38	6	F	30	41
39	6	F	27	27
40	6	F	25	27
Total:			824	1076

Table S2: (Related to Figure 2) The number of cells recorded in dmPFC and CA1 for each session.

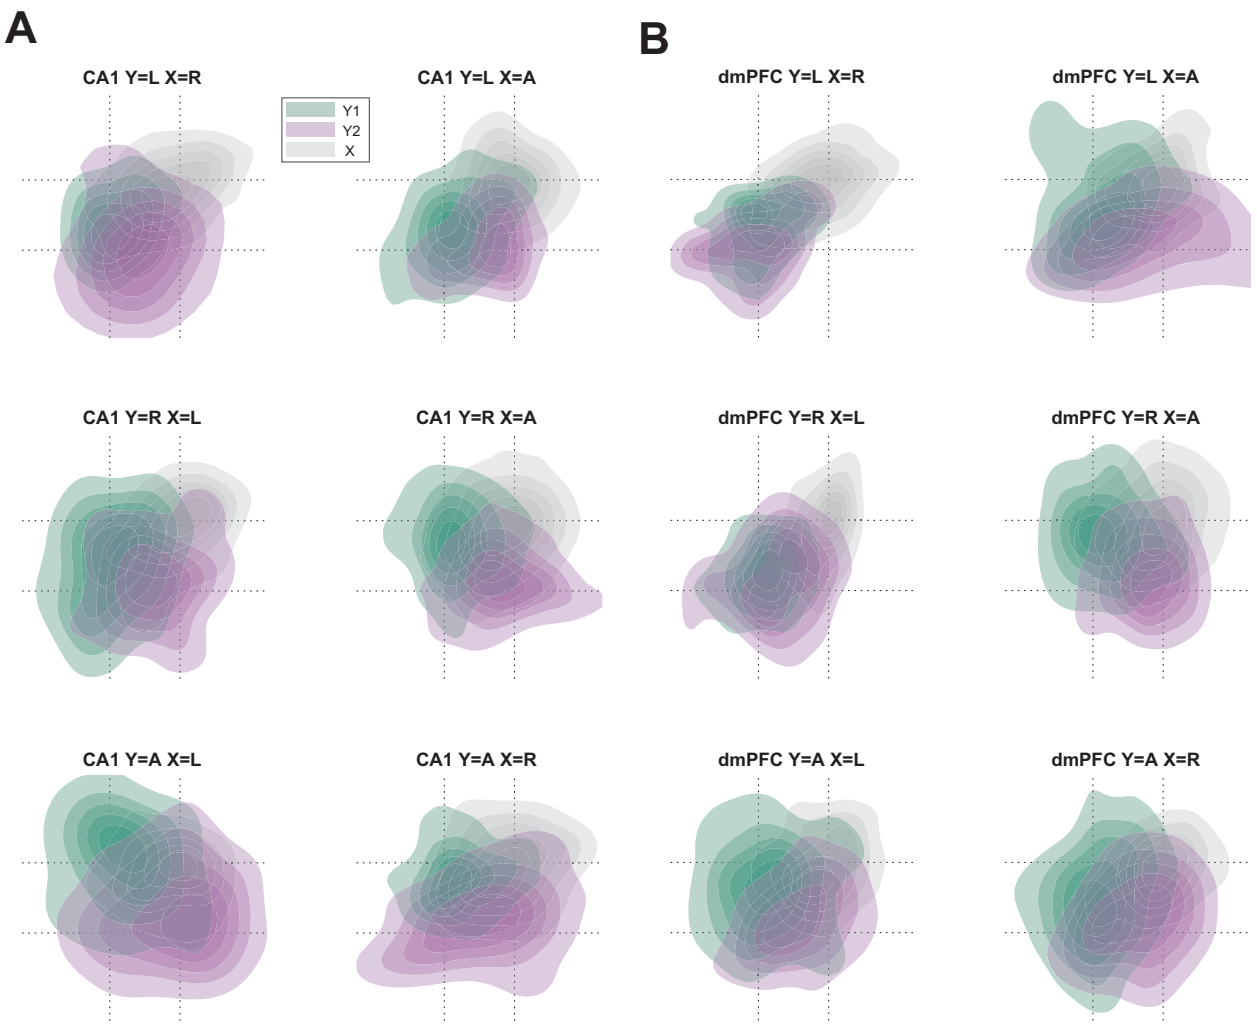

Figure S7: (Related to Figure 2) LDA projections in (A) CA1 and (B) dmPFC (as in Figure 2) for each transition type separately, for contingency triplets with only one intervening block. Y indicates the contingency type of the first and recurring contingency block, and X indicates the intervening block type.

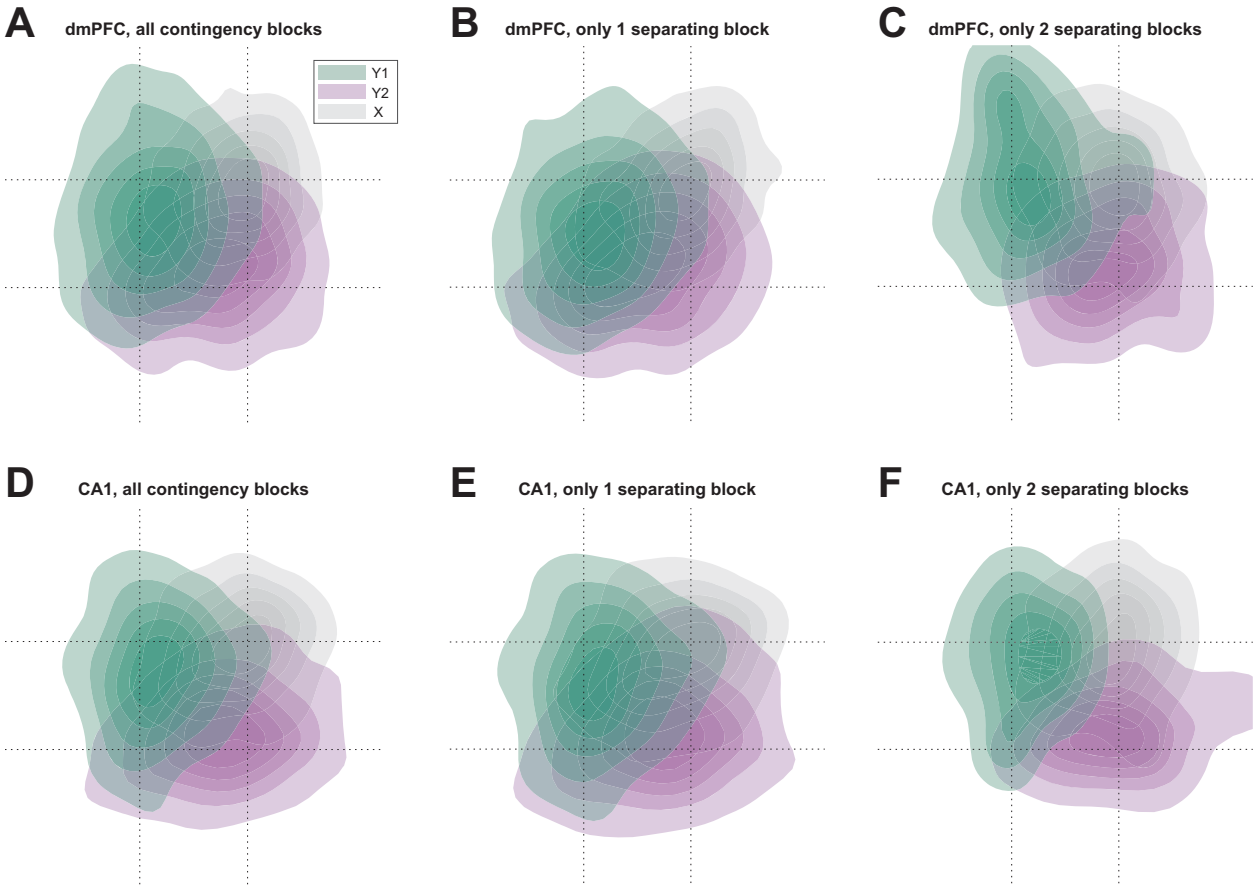

Figure S8: (Related to Figure 2) LDA projections for (A-C) dmPFC and (D-F) CA1 (as in Figure 2) for (A and D) all blocks, (B and E) for only $Y_1/X/Y_2$ contingency block sequences where X contained a single contingency block, and (E and F) for only $Y_1/X/Y_2$ contingency block sequences where X contained two contingency blocks.

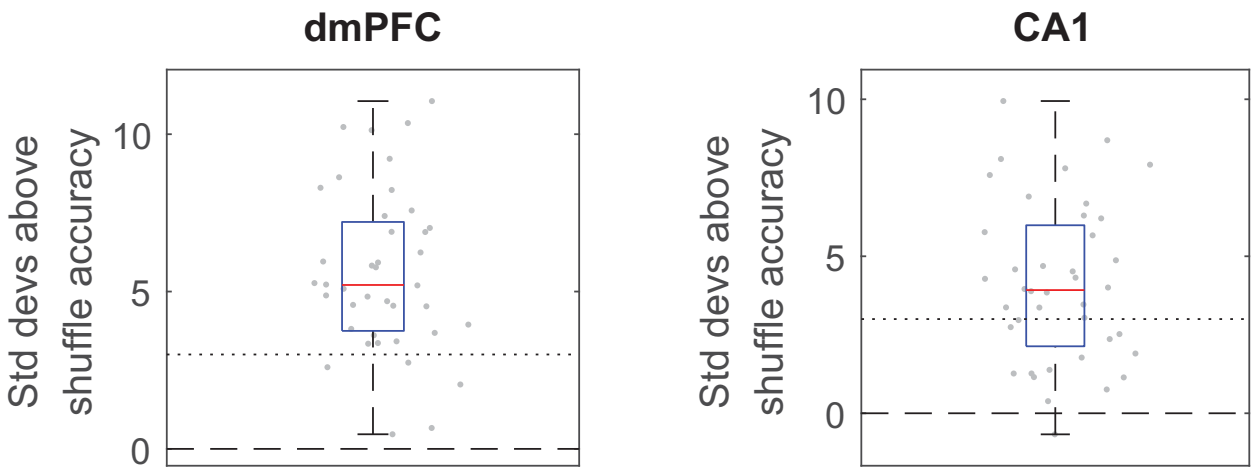

Figure S9: (Related to Figure 3) Contingency decoding from ensemble activity in dmPFC and CA1. Standard deviations above accuracy distributions resulting from shuffled spiketimes.

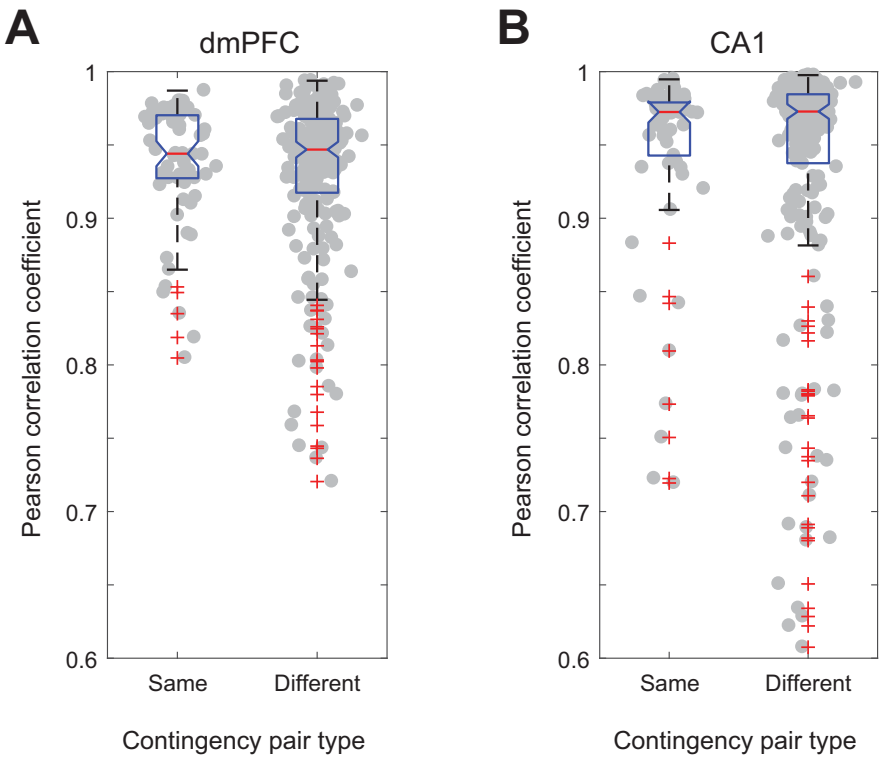

Figure S10: (Related to Figure 3) Ensembles do not appear to overtly undergo global remapping between contingency types. Pearson correlation between spatial tuning curves during passes through the maze's central track of ensembles in (A) dmPFC and (B) CA1.

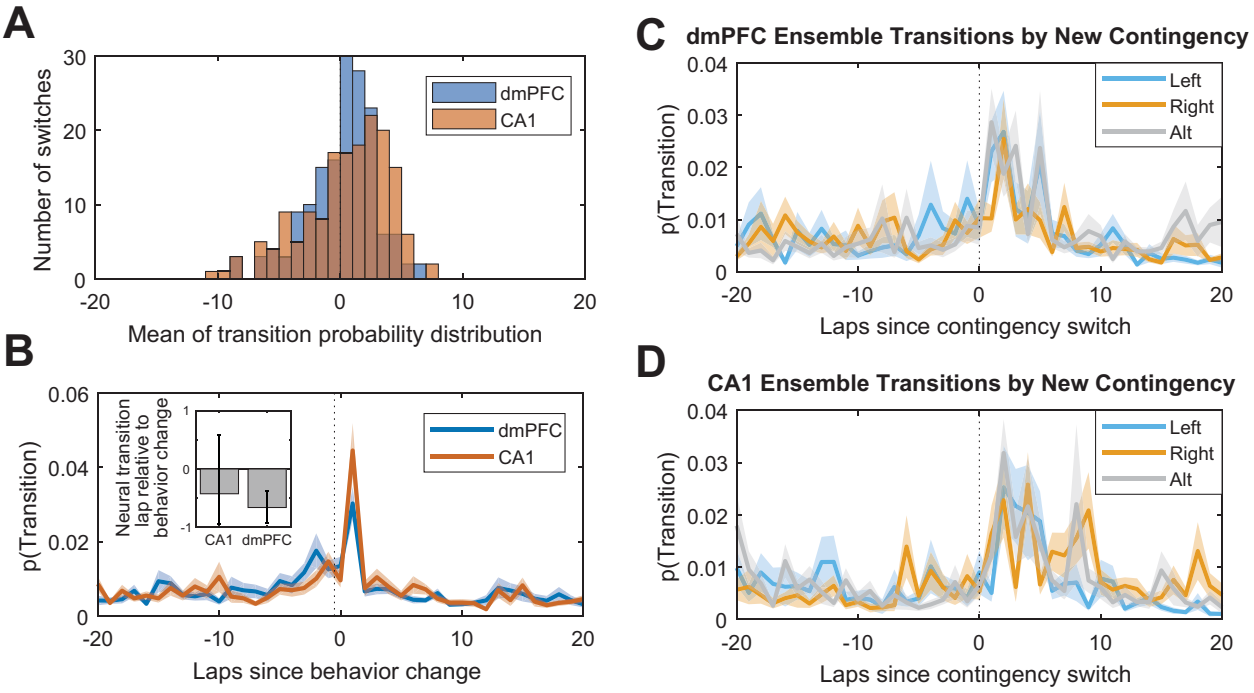

Figure S11: (Related to Figure 4) The timing of representational transitions in dmPFC and CA1. (A) Means of ensemble transition probability distributions relative to contingency switches for dmPFC and CA1. (B) Transition probability aligned to the behavior change. Shown is mean $\pm$ SEM, $N = 164$ contingency switches. Inset shows the mean neural transition lap relative to the behavioral change, with bootstrapped 95% confidence intervals. $N = 164$ contingency switches. (C) Transition probability split by new contingency in dmPFC. (D) Transition probability split by new contingency in CA1.

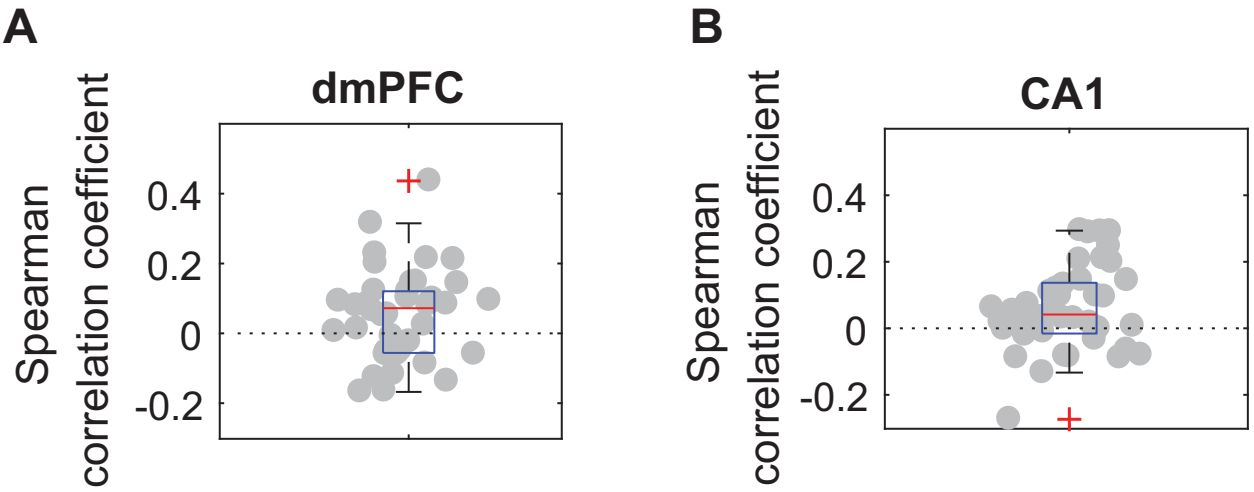

Figure S12: (Related to Figure 4) The correlation between LogIdPhi (a measure of the amount of VTE) and the probability of neural transitions in both (A) dmPFC and (B) CA1. $N = 40$ sessions.

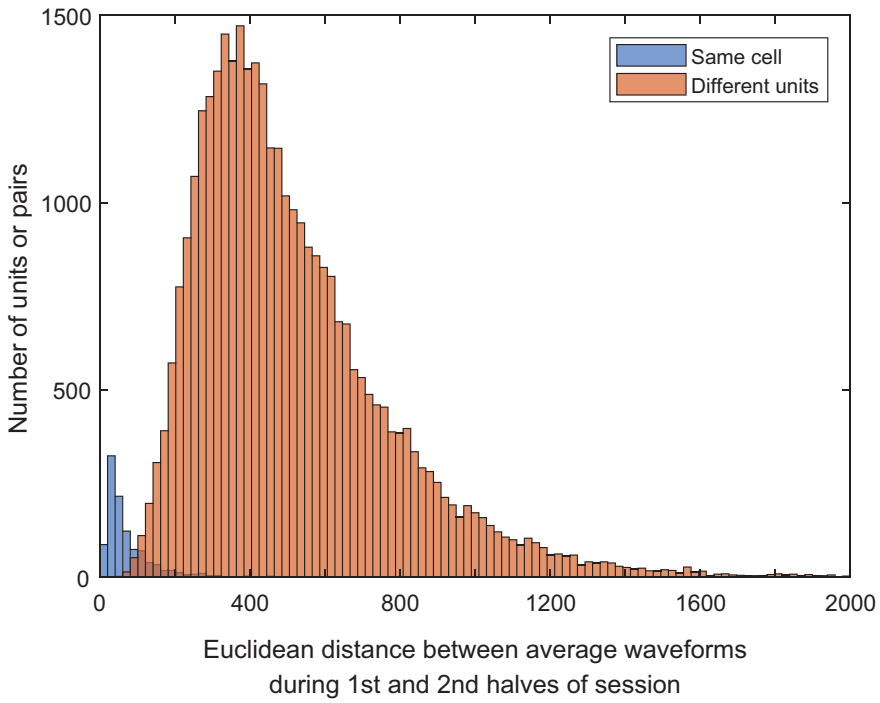

Figure S13: Distance of average spike waveforms from identified single units (blue) between the first half and second half of the session during which that unit was recorded. Also, the average spike waveform distance of unit pairs which were identified as different units during the same session (orange). Identified units had waveforms which were much more self-similar than to control unit waveforms, suggesting minimal recording drift. This analysis is similar to, but not identical to, the analysis performed in Tolias et al. (2007) to identify cell stability across days.