Published in final edited form as:
Neuron. 2017 April 05; 94(1): 183–192.e8. doi:10.1016/j.neuron.2017.03.011.

Medial Prefrontal Cortex Reduces Memory Interference by Modifying Hippocampal Encoding

Kevin G. Guise[1,2] and **Matthew L. Shapiro**[1,2,*]

[1]Department of Neuroscience, Ichan School of Medicine at Mount Sinai, 1 Gustave L. Levy Place, New York, NY 10029

[2]Friedman Brain Institute, Ichan School of Medicine at Mount Sinai, 1 Gustave L. Levy Place, New York, NY 10029

Summary

The prefrontal cortex (PFC) is crucial for accurate memory performance when prior knowledge interferes with new learning, but the mechanisms that minimize proactive interference are unknown. To investigate these, we assessed the influence of medial PFC (mPFC) activity on spatial learning and hippocampal coding in a plus maze task that requires both structures. mPFC inactivation did not impair spatial learning or retrieval per se, but impaired the ability to follow changing spatial rules. mPFC and CA1 ensembles recorded simultaneously predicted goal choices and tracked changing rules; inactivating mPFC attenuated CA1 prospective coding. mPFC activity modified CA1 codes during learning, which in turn predicted how quickly rats adapted to subsequent rule changes. The results suggest that task rules signaled by the mPFC become incorporated into hippocampal representations and support prospective coding. By this mechanism, mPFC activity prevents interference by "teaching" the hippocampus to retrieve distinct representations of similar circumstances.

Introduction

Memory impairs learning if new and established information conflict (Bartlett, 1932; Harlow, 1949) or when familiar stimuli must be associated with different responses (Greenberg and Underwood, 1950; Underwood, 1957), a phenomenon described as proactive interference (Bergström, 1893). Episodic memory represents stimuli in spatial, temporal, and personal context (Squire, 2004; Tulving, 1984), superordinate features that can reduce interference by dissociating stimulus identity and behavioral significance. For example, we know to answer a ringing telephone in our own homes, but not in others'. The hippocampus is crucial for episodic memory in people and other animals (Rempel-Clower et

---

[*] Lead Contact: matthew.shapiro@mssm.edu.

Author Contributions:
KGG and MLS designed the experiment, KGG carried out the experiments and analyses, and KGG and MLS wrote the manuscript.

al., 1996; Squire, 2004), and coping with proactive interference requires the prefrontal cortex. While the homologies between human and rodent PFC are debated (Preuss, 1995), proactive interference increases with PFC dysfunction in both species (Peters et al., 2013; Xu and Sudhof, 2013), and the neuronal mechanisms of this process are unknown.

Models of executive function suggest that the PFC resolves interference by modulating the relative activation strengths of memories during retrieval (Depue, 2012); it may also prevent interference by providing internal signals, e.g. task rules, that differentiate the significance of stimuli in the same external context. If so, then PFC signals should “pattern separate” otherwise overlapping hippocampal representations.

To investigate how the mPFC helps resolve memory interference, we tested the effects of mPFC inactivation in rats performing a hippocampus-dependent spatial reversal learning task. Though mPFC inactivation did not impair learning one spatial rule, it impaired all subsequent reversal learning, as though intact mPFC function differentiated rule-defined task episodes and sped subsequent learning by reducing proactive interference. To investigate the mechanisms by which functional interactions between the mPFC and the hippocampus might resolve memory interference, we recorded unit activity simultaneously in the mPFC and CA1 as rats performed the same task. We found that both structures discriminated task rules, that the mPFC influenced hippocampal activity, that mPFC inactivation reduced the separation of rule representations in the hippocampus, and that the magnitude of this effect predicted subsequent learning speed. By modulating hippocampal encoding, rule-related mPFC signals differentiate CA1 representations and reduce proactive interference in memory.

Results

To determine if mPFC activity altered spatial learning or memory, rats were trained on an elevated plus maze to perform a spatial memory task that requires hippocampal function (Ferbinteanu and Shapiro, 2003) (Figure 1A). In each daily testing session, a rat was placed on one of two potential start arms (north or south) and learned by trial and error to apply one of two mutually exclusive rules (‘go east’ or ‘go west’) to obtain food reward. After the rat made 10 correct responses in a block of 12 consecutive trials, the reward contingencies were reversed and the animal learned to use the opposite rule. After rats performed with ≥ 80% accuracy, they were tested with and without mPFC inactivation induced by bilateral infusions of either muscimol or vehicle 30 min prior to testing (Figure 1B). mPFC inactivation did not affect the initial discrimination (ID), but impaired subsequent reversal learning (Figure 1C; repeated measures ANOVA, drug × reversal interaction: $F(2,7) = 38.84$, $p < 0.001$; muscimol vs. saline post-hoc Bonferroni corrected t-tests: initial association: $t(9) = 2.33$, $p > 0.05$; first reversal: $t(9) = 8.79$, $p < 0.05$; second reversal: $t(8) = 11.39$, $p < 0.05$). The mPFC was not required for spatial learning or memory retrieval, but it was crucial for switching between rapidly changing spatial goals, consistent with findings that mPFC function reduces proactive interference among recent and highly overlapping events (Jonides and Nee, 2006; Shimamura et al., 1995).

mPFC activity pre-empts interference

mPFC activity could prevent proactive interference by detecting competing representations and modulating their relative activation during memory retrieval (Depue, 2012). For example, as a rat approaches the choice point, active memories of the previous discrimination could interfere with learning the new competing rule, and mPFC signals could suppress the incorrect or enhance the correct one. Alternatively, the mPFC could mitigate interference during learning by modulating encoding, e.g. by integrating task rules (Miller and Cohen, 2001) that differentiate spatial representations. Because the mPFC was inactivated throughout testing first experiment did not address this question.

The next experiment distinguished between these possibilities by comparing the effects of inactivating the mPFC either before or after the rats learned the ID. The new procedure included both a sham and an actual infusion and separated the ID and the first reversal (R1) by a 30 min delay (Figure 1D). If the mPFC modulates active memories during retrieval, then inactivating the mPFC *after* the ID and during memory retrieval should impair R1. If the mPFC modulates memory acquisition, however, then inactivating the mPFC after the ID should not impair R1 because the first rule (e.g., 'go east') would have been integrated normally with hippocampal representations. Subsequent reversals (R2, R3...) should be impaired, however, because the representations would be encoded without the mPFC rule signals that distinguish these otherwise ambiguous situations.

As in the first experiment, inactivating the mPFC before the ID impaired R1 but not ID, and the additional delay and sham infusion had no effect on performance (Figure 1E; drug × reversal interaction: $F(1,7) = 62.07$, $p < 0.001$; post-hoc t-tests between drug conditions: ID: $t(7) = 0.73$, $p > 0.05$; R1: $t(7) = 10.58$, $p < 0.001$; R2: $t(3) = 6.61$, $p < 0.05$, cf. Figure 1C). Switching the order of the actual and sham infusions inactivated the mPFC *after* the ID and *before* reversal learning had no effect on R1, demonstrating that mPFC activity was not needed to switch from one goal to another. Subsequent reversals were impaired, however, indicating that mPFC activity was needed during learning if a new goal was to be learned subsequently (drug × reversal interaction: $F(2,14) = 14.17$, $p < 0.001$; ID post-hoc: $t(7) = 0.16$, $p > 0.05$; R1 post-hoc: $t(7) = 1.06$, $p > 0.05$; R2 post-hoc: $t(7) = 4.91$, $p < 0.01$). The combined results show that mPFC activity supported reversals during learning, rather than during memory retrieval, and suggest that mPFC function may prevent proactive interference by differentiating memory representations during encoding.

To investigate if mPFC modulated hippocampal codes, we recorded mPFC and CA1 activity simultaneously while animals performed the spatial reversal task and analyzed local field potentials (LFPs) and single unit ensembles during learning. This paper focuses on activity in the start arm, where identical spatial behaviors precede distinct choices, and when memory could inform goal selection. For these analyses, we defined task rules operationally as the goal arm entered by the animal, and investigated 'rule coding' by comparing neuronal activity in 'east-going' and 'west-going' trials, as rats used identical behaviors in the same start arm to approach the choice point(Ferbinteanu and Shapiro, 2003; Rainer et al., 1999). Single units in both structures fired at different rates in the start arms on the way to different goals, thereby showing prospective coding (examples in Figure S3.2).

mPFC and CA1 activity is synchronized through theta frequency oscillations—If the mPFC minimizes interference by modulating hippocampal coding, then activity in the two structures should be temporally coordinated. Theta (4–12 Hz) and gamma (30–80 Hz) oscillations are prominent features of hippocampal activity and are believed to provide the synchronization needed for efficient transfer of information between the hippocampus and other parts of the brain (Spellman et al., 2015). Each oscillation frequency is observed in the mPFC and modulated by the phase of hippocampal theta during memory-guided behavior (Lansink et al., 2016; Sirota et al., 2008).

To test how oscillations in the mPFC and CA1 were coordinated during the spatial reversal task, we analyzed LFPs and spiking activity from the time the rat was placed in a start arm facing away from the maze center (time −1 s) and started moving along the arm (time 0 s, distance = 0), through the choice point (time ~1s, distance = 60 cm) until it reached the reward point ~2.5 s later (distance = 120 cm, Figure 2A). LFP spectra revealed prominent theta oscillations in CA1; theta power was highest as animals traversed the start arm and declined as the animal entered the goal arm (Figure 2B). In contrast, the pattern of oscillations in the mPFC of typical trials started with low-frequency, high-amplitude oscillations in the 4 Hz range, as previously reported (Fujisawa and Buzsaki, 2011). As animals initiated forward movement on the start arm, mPFC oscillatory power decreased while the frequency increased to about 5 Hz. Theta oscillations became prominent in mPFC as animals moved into the goal arm, and the oscillation frequency increased to about 10–12 Hz as rats approached the reward point. Theta oscillations in the mPFC and CA1 were coherent in all arms of the plus maze, with the strongest coherence observed before animals initiated a trial and in the goal arm.

To test if coherence reflected coordinated neuronal activity (Kajikawa and Schroeder, 2011), we examined the extent to which unit activity in the prefrontal cortex was phase modulated by the ongoing hippocampal theta rhythm. Indeed, both CA1 units and mPFC units were phase locked to hippocampal theta: 75.8% (235/310) of CA1 units were significantly phase modulated, as were 20.2% (45/223) of mPFC units ($p < 0.05$, FDR corrected (Benjamini, 1995); Figure 2C). mPFC units tended to fire ~ 47 ms before the theta trough, whereas CA1 units fired ~4 ms after the trough, corresponding to a temporal offset with mPFC units either leading CA1 units by about 51 ms or lagging by 67 ms (Figure 2D). Before rats entered the choice point, mPFC and CA1 oscillations were synchronized on a timescale that could permit activity in each structure to influence the other (Figure 2E).

mPFC and CA1 population activity track task rules—To assess ensemble activity during learning we represented each trial as a rate vector based on the firing rate of each neuron (CA1 N = 14.82 / session, s.d. = 6.8; mPFC N = 11.8 / session, s.d. = 4.3) in a restricted part of the start arm. Vectors were first divided into seven equally spaced bins, from the first to the final trials of each contingency, and then averaged within bins. Ensemble activity dynamics were quantified by calculating the similarity between average rate vectors among all pairs of bins (see supplemental information). To visualize learning-related changes in mPFC and CA1 activity, we used multidimensional scaling (MDS) to generate two-dimensional projections of a 3D space that represented the similarity of rate vectors across all trials (Figures 3A, B, supplemental information). Every point in each MDS plot corresponds to the activity state in one brain region during a single epoch; the proximity of points shows the similarity of population vectors, and each triangle indicates activity at the seven points across the learning curve. The MDS plots reveal dynamic activity states in mPFC and CA1 that correspond to learning each task rule (i.e., ‘go east’ or ‘go west’), with distinct activity patterns separating the ID and each reversal. During stable performance within a contingency, brain activity was stable, and adjacent trials clustered in the same region of the activity space (e.g. the triangles indicating Goal #2 in the lower right of Figure 3A). When contingencies changed, so did ensemble activity, illustrated by trajectories through the activity space from a region corresponding to one goal toward another goal. CA1 and mPFC ensemble activity in the start arm appeared to distinguish task rules along the vertical axis and time along the horizontal axis. The firing rate patterns of a subset of individual units discriminated task rules via prospective coding (Figure S3.2).

To quantify the extent to which ensemble codes predicted the rats’ choices, we used support vector machines (SVMs) to categorize the task rule represented in the start arm of all single trials, and compared these to SVMs trained on the same data with shuffled goal labels using permutation tests. CA1 ensembles predicted the pending choice of each trial with a mean accuracy of 82.3%, and 24 of 28 ensembles predicted the goal better than 95% of the permutations (Figure 3C; expected (5%): 1.4 out of 28; χ^2 (1) = 364.83, $p < 0.05$). mPFC ensembles predicted the goal of single trials with a mean accuracy of 74.0%, and 20 out of 28 ensembles predicted the goal significantly better than chance (expected (5%): 1.4 out of 28; χ^2 (1) =247.11, $p < 0.05$). Because mPFC and CA1 ensemble activity distinguished the goal of each trial well before the discriminative response, each structure could inform downstream brain circuits during memory retrieval.

Ensemble firing rates predict choices, but do not represent rules or goals— mPFC and CA1 ensembles could inform choices by representing task rules, e.g., “go East”, spatial goals, e.g., an “image” of the correct goal arm, a pending egocentric response, e.g., “turn left”, or more abstract signals that communicate distinct messages (Shannon, 1948). To investigate if prospective activity represented task features, we compared ensemble activity in pairs of arms that might share commonalities. We reasoned that if neural activity represented “goal images”, then ensembles should be similar in the start and selected goal arm, whereas if activity represented abstract rules, then ensembles in the two start arms should be correlated on journeys to the same goal. We therefore quantified “goal representation” as population vector correlations (PVRs) between each start arm and the selected goal arm, “rule representation” as the PVRs between the two start arms in journeys to the same goal, and “egocentric representation” as the PVRs between the two start arms in journeys to opposite goals. By these measures prospective activity did not represent goals, rules, or body turns, but strongly differentiated each type of pending journey. Rather than predicting the pending choice with overlapping “goal image” representations of the start and goal arms, PVs were weakly anticorrelated in both structures (PVR(start, goal): mPFC: $r = -0.06$, $t(27) = 5.8$, $p < 0.001$; CA1: $r = -0.11$, $t(27) = 5.36$, $p < 0.001$). Similarly, PVs in the two start arms were either uncorrelated or anticorrelated in journeys to the same goal rather than representing the common “rule” with overlapping activity patterns (r(N,S): mPFC: $r = 0.002$, $t(27) = 0.17$, $p = 0.86$; CA1: $r = -0.22$, $t(27) = 12.51$, $p < 0.001$). Journeys from each start arm require opposite body turns to enter the same goal. If PVs in the two start arms represent pending egocentric responses, then they should be positively correlated when the same body turn leads to opposite goals. Again, PVs in the two start arm were either uncorrelated or anticorrelated prior to the egocentric response, ($r(N \rightarrow E, S \rightarrow W)$: mPFC $r = -0.006$, $t(27) = 0.51$, $p = 0.35$; CA1 $r = -0.19$, $t(27) = 6.6621$, $p < 0.001$). Moreover, the correlation between start arm PVs did not distinguish identical and opposite body turns (mPFC: $t(27) = 0.89$, $p = 0.38$; CA1: $t(27) = 1.35$, $p = 0.19$). Together, the PVRs show that neither mPFC nor CA1 prospective codes represent rules, goals, or egocentric patterns in different journeys. The results suggest instead that prospective codes differentiate episodes guided by similar task features. Because we analyzed only ensemble firing rates, however, the results do not exclude rule, goal, or any other representation coded by different mechanisms, e.g., firing sequences related to vicarious trial and error (Johnson and Redish, 2007).

Changes in mPFC activity predict changes in CA1 activity during learning— mPFC and CA1 ensembles could predict goals independently or cooperatively, simultaneously or sequentially. We hypothesized that rule-related mPFC activity modulates CA1 encoding, so that changes in mPFC activity should predict CA1 dynamics as learning proceeds. Alternatively, if place-related CA1 activity helps select mPFC codes, then changes in CA1 activity should predict those in mPFC. We compared temporal relationships between mPFC and CA1 ensemble activity in the start arm using Granger prediction analysis, which measures the extent to which one time series predicts changes in another beyond what each series predicts about itself (Cohen, 2014). We assessed the three trials that immediately followed a rule change, and used the activity in the first two trials to predict activity in the third. We found that the recent history of mPFC ensemble activity predicted changes in CA1 ensembles (Figure 3D, see supplemental information for methods). mPFC activity during the previous two trials increased the explained variance in CA1 activity during the next trial by 6.4% ($p < 0.001$, permutation test). In contrast, CA1 ensemble activity did not predict mPFC activity during the same trials (0.8% increase, $p > 0.05$, permutation test). During stable performance, however, the predictive relationship reversed and declined, and CA1 activity reliably predicted mPFC activity (< 1% increase in variance explained, $p < 0.05$, permutation test). The results show that predictive coding is asymmetric and emerges sequentially, not simultaneously, in mPFC and CA1 during learning, and suggest that mPFC activity modulates the formation of CA1 representations. The results so far do not prove causality, however: earlier changes in mPFC compared to CA1 activity could reflect, e.g., different circuit responses to a common input.

mPFC activity differentiates CA1 prospective coding— If mPFC indeed modulates CA1, then mPFC inactivation should reduce predictive coding by CA1. To test this inference we inactivated the mPFC unilaterally and recorded ipsilateral CA1 activity. Because unilateral mPFC inactivation had no measurable effect on observed behavior or reversal learning (Figure S2.1), its effects on CA1 activity was unlikely due to altered behavior. If the mPFC reduces interference by enhancing the distinction between CA1 representations, then mPFC inactivation should reduce CA1 prospective coding separation, measured as the difference between ensemble activity states when rats are guided by different goals.

Ipsilateral mPFC inactivation did not reduce overall choice prediction accuracy by CA1 ensembles (SVM decoded goals: saline: 79.2%; muscimol: 79.1%; Ensemble size: saline: N = 15.2 / session, s.d. = 6.7; muscimol: N = 15.3 / session, s.d. = 8.3), and most CA1 ensembles continued to classify goals better than chance (saline: 10 out of 14, 71.6%, expected (5%): 0.70, χ^2 (1) = 123.56, $p < 0.05$; muscimol: 11 out of 12, 91.7%, expected (5%): 0.6, χ^2 (1) = 180.27, $p < 0.05$; Figure 4A). At the same time, mPFC inactivation reduced the separation between CA1 prospective codes, quantified as the distance between activity in each trial and the SVM classification margin, the boundary separating pending East and West choice representations. CA1 prospective codes were less distinct when the mPFC was inactivated (Figure 4B) (mean per-session z-scored margin width: saline = 0.854, s.d. = 0.26; muscimol= 0.665, s.d. = 0.20, permutation tests, $p < 0.05$), demonstrating that unilateral mPFC inactivation reduced pattern separation in ipsilateral CA1 representations.

mPFC inactivation does not affect CA1 place coding—mPFC inactivation reduced predictive prospective code separation in CA1 ensembles, implying that mPFC activity modulated CA1 during learning. The results so far do not determine if mPFC activity influenced all hippocampal computations or a subset related to task demands. Hippocampal neurons fire in place fields that reflect spatial computations and guide navigation (O'Keefe, 1971), and mPFC activity could modulate these computations or, e.g. attention (Cassaday et al., 2014). If mPFC activity modulates CA1 computations generally, then mPFC inactivation should reduce spatial pattern separation as well as predictive goal coding. We therefore measured place discrimination by CA1 neurons using the same SVM methods described earlier, and found that CA1 ensembles discriminated the two start arms on each trial with high accuracy whether or not the mPFC was inactivated (mean decoding accuracy, saline: 98.9%; muscimol 99.3%). Every CA1 ensemble supported the discrimination (saline: 14 of 14, expected (5%) = 0.70, χ^2 (1) = 252.7, $p < 0.05$; muscimol: 12 of 12, expected (5%) = 0.6, χ^2 (1) = 216.6, $p < 0.05$; Figure 4C), and mPFC inactivation did not affect the magnitude of place discrimination (SVM margin width, saline: mean = 1.79, s.d. = 1.32; muscimol: mean = 1.97, s.d. = 1.19; $p > 0.05$; Figure 4D). Established measures of place coding extended the SVM findings: mPFC inactivation resulted in a small, non-significant reduction in CA1 spatial information content (c.f.,(Kyd and Bilkey, 2003)). The spatial information per spike for CA1 units in the saline condition was 2.37 (s.d. = 1.77), and 2.35 (s.d. = 1.83) for muscimol (difference: $p > 0.05$, permutation test). mPFC inactivation affected goal but not place coding, suggesting that different networks maintain spatial representations as mPFC modulates CA1 rule coding.

mPFC modulation of CA1 predicts subsequent learning rate—The results so far suggest that mPFC rule signals modify CA1 encoding, and prevent interference by promoting the formation of distinct memory representations. If mPFC separation of CA1 prospective codes improves subsequent learning, then future performance should vary with mPFC's current influence. In quantitative terms, the magnitude of the Granger value measured as one rule is learned should predict how quickly the next rule is acquired. We therefore compared the Granger prediction in sets of trials that preceded reversals learned at different rates, categorized as fast, moderate, or slow learning speeds (1–3, 4–7, or >7 trials to above chance performance; (Smith et al., 2004)). The magnitude of mPFC's effect on

CA1 activity predicted the speed of subsequent rule learning. The largest effect of mPFC on CA1 dynamics occurred during early rule learning and stable performance in trials that preceded the most rapid reversals, when the next rule was learned in <= 3 trials (increase in explained variance: early learning = 29.0%, stable performance = 4.8%, permutation tests $p < 0.05$, Figure 5). mPFC had less effect on CA1 dynamics when subsequent reversal learning proceeded in > 3 trials. Moreover, the predictions pertained only to subsequent learning. mPFC effects on CA1 dynamics were weak when learning speeds were categorized using either the previous or the current rule (maximum increase in explained CA1 variance by mPFC: previous rule < 7.1%, $p < 0.05$; current rule, < 8.3%, $p < 0.05$). Though both mPFC and CA1 ensemble activity predicted goal choices, switching to a new rule was fastest when mPFC activity most strongly modulated CA1 as rats learned the previous rule. The mPFC may support reversal learning by increasing goal-separated prospective codes in CA1 that reduce interference proactively across episodes.

Discussion

In a hippocampus-dependent spatial task, mPFC dysfunction left discrimination learning and retrieval intact but impaired switching between rules in serial reversals. mPFC and CA1 spiking was theta synchronized, ensemble activity in each structure predicted pending choices, mPFC inactivation reduced prospective, but not place coding in CA1 representations, and mPFC modulation of CA1 activity as rats learned one rule predicted how quickly they switched to the next rule. These findings are consistent with the view that mPFC computations improve hippocampal-dependent learning and memory by signaling changing goals, rules that distinguish representations of identical places in the same spatial context.

Encoding, retrieval, and memory interference

The mPFC utilizes contextual information to minimize interference, providing task rules that discriminate amongst conflicting representations (Preston and Eichenbaum, 2013). While the hippocampus is ideally situated anatomically to associate multiple aspects of experience that overlap in place and time, the mPFC further integrates representations from the hippocampus (Preston and Eichenbaum, 2013) and other regions (Hoover and Vertes, 2007). mPFC circuits may help compute neural representations of abstract rules, superordinate representations of different behaviors, e.g. “turn left” and “turn right”, performed to accomplish the same goals, e.g. ‘go east’ or ‘go west.’ Rules, motives, and other goal-related signals define internal contexts that, like space and time, can integrate or segregate memories of past and present experience.

The current findings extend the role of the mPFC in hippocampal-dependent learning and memory from coping with interference from the past to minimizing potential interference in the future, and imply that mPFC inputs contribute to learning and memory retrieval through Hebbian mechanisms. As input to the hippocampus, different internal context signals could reduce the overlap among active representations when other conditions are similar. If distinct input patterns modify synaptic weights during learning, then their subsequent activity will help activate rule-separated hippocampal representations, thereby reducing interference during memory retrieval. This view helps integrate disparate studies by proposing that mPFC activity during learning is coupled to its subsequent role in memory retrieval.

The present experiment trained rats to switch between spatial goals in a fixed external context, so that appropriate responses relied on internal context and had to be learned in each reversal. In these conditions, mPFC inactivation impaired neither retrieval of an established spatial memory nor learning one new spatial goal. Rather, performance was impaired only if the mPFC was inactivated during both initial and reversal learning (Figure 1E), when identical stimuli were associated with conflicting responses, as though mPFC activity during encoding was necessary for its subsequent contribution to retrieval.

mPFC inactivation impairs contextual memory retrieval in tasks that require the hippocampus (Navawongse and Eichenbaum, 2013; Peters et al., 2013; Xu and Sudhof, 2013). For example, rats were presented with both lemon and peppermint scented cups in two spatial environments, and trained to associate one odor with food in each. In this case, intact rats learned two fixed associations between external context, stimulus, and reward, and mPFC inactivation after training impaired memory retrieval (Navawongse and Eichenbaum, 2013). If the mPFC is needed to learn contextual associations, then inactivating the mPFC during training should prevent the formation of context-separated odor-reward associations and thereby preclude subsequent contextual retrieval.

Prospective coding by CA1 ensembles differentiates imminent goal choices, likely contributes to memory retrieval, and may require coordinated activity in the hippocampus and PFC. Spatial memory tasks favoring a prospective strategy are impaired by both bilateral mPFC inactivation as well as crossed inactivation of the mPFC and the contralateral hippocampus (Goto and Grace, 2008). Recording studies concur that mPFC activity contributes to prospective coding by hippocampal ensembles, and suggest that rule-based pattern separation supports contextual retrieval. In another spatial task, prospective coding by hippocampal neurons distinguished pending trajectories (e.g., go left vs. go right), and inactivating the nucleus reuniens, which relays signals from the mPFC to the hippocampus, eliminated the prospective code (Ito et al., 2015). Unilateral inactivation of the mPFC reduced, but did not eliminate prospective coding separation in the present experiment, perhaps because other inputs from the PFC remained intact, e.g. the OFC and anterior cingulate that would be disconnected from the hippocampus by nucleus reuniens inactivation (McKenna and Vertes, 2004; Vertes et al., 2006). CA1 activity distinguished between identical places and scents in different spatial contexts, and though place fields were maintained after mPFC inactivation, firing patterns no longer discriminated the objects during stimulus sampling, as if rule information was lost (Navawongse and Eichenbaum, 2013). The present results agree that mPFC activity modulates CA1 prospective coding, and show that the strength of this modulation is strongest during reversal learning, as hippocampus ensembles form predictive representations. The data suggest that CA1 ensembles “learn to retrieve” by incorporating mPFC “rule” signals that subsequently activate context-separated representations. The results show that CA1 prospective coding differentiates rules and egocentric responses as well as spatial goals, suggesting it contributes generally to memory discrimination, distinguishing similar events by their association with other salient task features. By linking internal or external contextual features to ambiguous stimuli, mPFC signals establish separate CA1 representations of similar events during learning and activate prospective codes during retrieval that reduce memory interference.

While the PFC may reduce interference during both encoding and retrieval, its contribution to behavior likely depends on task demands that interact with a limited functional capacity. In people, proactive strategies decline as demand for “real time” working memory increases (Braver, 2012), as though keeping current and potential future items “in mind” require the same computational resources. As in people, mPFC activity in rats may guide present or future behavior as ongoing tasks demands vary. For example, increasing the frequency of spatial reversal discriminations (Riceberg and Shapiro, 2012) or the number of simultaneous discriminations (Peters et al., 2013) may “shift” mPFC resources away from memory encoding and toward handling interference among co-active representations. The mechanisms affected by task demands that shift PFC resources are unknown, and may depend on how discriminative information is presented in time. Discrete trials separated by seconds to minutes include behavioral pauses typically accompanied by hippocampal sharp wave ripples (SWRs) that coordinate PFC and CA1 activity and support memory consolidation (Jadhav et al., 2016), conditions that may promote mPFC integration with hippocampal encoding. Continuous trials separated by shorter delays may reduce behavioral pauses, SWRs, and opportunities for integrating mPFC signals with hippocampal encoding.

In summary, behavioral and physiological evidence suggest that mPFC activity during learning modifies hippocampal encoding to proactively minimize subsequent interference among representations during memory retrieval (Figure S4.1). mPFC activity conveys rules in the present that help separate hippocampal representations and facilitate subsequent learning; eliminating mPFC signals reduces rule discrimination in hippocampal codes. mPFC signals related to abstract rules exemplify one type of internal context that can help memory discriminations when the world appears stable but the situation has changed.

STAR Methods

CONTACT FOR REAGENT AND RESOURCE SHARING

Requests for additional information may be directed to, and will be fulfilled by, the lead contact, Dr. Matthew Shapiro (matthew.shapiro@mssm.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAIL

All experiments described in this paper used Adult Long-Evans rats purchased from Charles River Laboratories (N = 17; age: ~2 months; weight: ~280 grams). Rats acclimated to the colony for one week with ad-lib access to food and then started on food restriction. Food intake was monitored and adjusted so that the rats maintained at least 80% of their projected ad-lib weight, supplemented for normal growth. After 5 days of food restriction animals were handled by the experimenter for 20 minutes a day for 5 days to acclimate them to human contact.

METHOD DETAILS

Behavioral testing

Apparatus: Behavioral testing and recording was performed on an elevated plus maze (arms 64 cm long, 6.5 cm wide) and a waiting platform were centered in a room decorated with high-contrast posters to allow the rats to visually orient themselves in space. The maze was painted matte black and was fitted with infrared emitter/sensor pairs to automatically log animals' responses and to trigger LED cue onsets. The emitter/sensor electronics were purpose built, interfaced with a PC via a NI-DAQ 6008 (National Instruments), and controlled using in-house software written in Matlab.

Behavior Training

Shaping: Animals were introduced to the behavioral testing room after handling. The experimenter placed chocolate sprinkles in the ends of the goal arms of the plus maze and on the wait platform, where he then placed a rat. After the rat consumed all the food on the platform, it was placed on the north start arm facing away from the choice point and allowed to explore the maze until all the chocolate sprinkles at the end of both of the goal arms were consumed. The rat was then returned to the wait platform, and the experimenter re-baited the goal arms of the maze. After ~1 minute, the rat was placed on the south start arm facing away from the choice point and allowed to explore the maze until the chocolate sprinkles in both goal arms were consumed. The rat was then returned to the wait platform for 1 minute and then returned to its home cage. Training on the behavioral tasks started on the following day

Training: One group of rats (N = 4) was trained only in the spatial reversal task. Another group (N = 6) was trained in both the spatial reversal task and a cue-response task (see next two headings). Training in both tasks occurred in separate sessions separated by approximately 6 hours, and the order of task presentation was randomized day to day. The behavioral results described in the main text of the paper were indistinguishable in the two groups of rats (Figure S1), and data from the cue-response task are not discussed here. After a rat reached the training criterion (see below), it was implanted with cannula in the mPFC bilaterally (Figure S4.2) and allowed to recover for 7 days before being handled again. Animals were then re-trained to criterion performance before behavioral testing resumed. A second group of animals was trained similarly and implanted either with dual site hyperdrives targeting mPFC and CA1 (N = 3), or with a cannula in the right mPFC and a recording hyperdrive targeting the ipsilateral CA1 (N=4).

Spatial reversal task: Rats were trained to perform a spatial memory task that requires hippocampal function. In each trial, the experimenter put food at the end of one goal arm in the plus maze, picked up the rat from the wait platform, and placed it on one of the two start arms (north or south); the rat was allowed to enter a goal arm, and was then returned to the wait platform. The same goal arm was rewarded throughout a block of trials, when the rat learned by trial and error to apply one of two mutually exclusive rules ('go east' or 'go west') to find the food. After a correct choice the rat was allowed to consume the reward before it was returned to the wait platform. Errors trials were counted when a rat put all four paws into the non-rewarded goal arm. During the first three trials of a block the animals were allowed to self-correct; otherwise, the error trial ended when the rat reached the end of the incorrect goal arm or turned around. After the rat made 10 correct responses in a block of 12 consecutive trials, the reward contingencies were reversed and the animal learned to use the opposite rule to find food. This procedure was repeated until animals completed ~64 trials. If a rat did not learn a spatial discrimination within 20 trials it was returned to its home cage. Animals were given rat chow in their home cages after >= 2 hours had elapsed since the end of testing on a given day. Rats were trained in the task until they performed at or above 80% accuracy for 3 out of 4 consecutive training days.

The start arm sequence was chosen pseudorandomly so that the rank correlation between subsequent start arms was less than +/− 0.100 and the same start arm was not repeated more than 4 times in four consecutive trials. The identity of the initial goal arm was chosen pseudorandomly across days.

Cue-Response Task: A group of animals (N=6) was trained to perform a cue-response task as well as the spatial reversal task, and the cue-response data are not described in the main text. As described in the supplemental material, animals trained in one or both tasks performed indistinguishably. Animals were trained on the cue-response task in the same testing room, but in near-complete darkness to minimize interference from spatial cues around the plus maze (deep-red light allowed the experimenter to see). All behavioral procedures were identical to the spatial task, except that the entrance to the correct goal arm was signaled by LEDs that were illuminated when the rat reached the halfway point of the start arm. Entering the cued goal arm led to the reward. The identity of the goal arm was chosen pseudorandomly in the same fashion as the start arms, with the additional criterion that the rank correlation between the start and goal arm sequences was less than +/− 0.100.

Cannula infusions: During the infusion procedure, animals were held in the experimenter’s lap and given a small piece of rat chow while the dummy cannula were removed and the infusion cannula were inserted into the guides. No anesthesia was used, and cannula were held in place with a small piece of masking tape. Animals were pet gently to keep them occupied during the infusion, and were not otherwise restrained. The infusion cannula were connected to 1 μL Hamilton syringes and an electronic pump delivered the infusate (muscimol or saline; see below) at a rate of 0.25 μL/min for two minutes (0.5 μL total). The infusion cannula were left in place for an additional five minutes to allow for diffusion of the drug/vehicle. When the infusion was complete, animals were given another piece of rat chow while the injection cannula were replaced with the dummy cannula. The animals were then returned to their home cages for 25 minutes prior to the start of testing.

Muscimol was dissolved in phosphate buffered saline (PBS) (0.1 μg/μL), and PBS was used as the control infusate (Rich and Shapiro, 2007; Young and Shapiro, 2009).

Delayed infusion experiments: A second set of experiments added a delay between ID and R1 to compare the effects of mPFC inactivation given before or after ID. The first experiment delivered infusions before the ID as described earlier. In the delayed infusion procedure, the animals were removed from the testing room after the ID, given a sham infusion that mimicked the sensory aspects of the actual infusion using empty cannula and infusing nothing, and then placed on the wait platform for 25 min before R1. The second experiment switched the order of the actual and sham infusions, so that the sham infusion preceded the ID and the actual infusion preceded R1.

Task order: The infusion/testing schedule followed an ABBA design, with saline infusions given on days 1 and 4, and muscimol on days 2 and 3. Behavioral testing started 25 minutes after the infusion. For the first four days of testing, six of the animals performed the spatial task first and the cue task immediately after. The other four animals performed only the spatial task. The animals trained in both tasks repeated the first four days of testing, except that they performed the cue task first and the spatial task second. During these days, infusions took place 50 minutes before the spatial reversal task, as opposed to the usual 25. Figure S1 illustrates that training on the additional task had no effect on the observed interference effect following mPFC inactivation.

Electrophysiology

Hyperdrive implants and tetrode lowering: Hyperdrives with 24 independently movable tetrodes were built in-house. Tetrodes were spun from 12.5 μm nichrome wire (Kanthal Precision Technologies), loaded into the hyperdrive, and cut and gold plated until the impedance on each wire was approximately 200 kΩ measured at 1000 Hz. During implantation, the electrode interface board (EIB 36 24TT; Neuralynx) was connected via 0.003” stainless steel wire to four ground screws distributed across the skull as well as two reference screws implanted above the cerebellum. The implant coordinates in mm from Bregma were CA1: AP −3.6, ML 2.0; mPFC: AP +3.0, ML 0.5. See the section on general surgical procedures for more detail. The tetrodes were lowered 1.4 mm into the cortex after surgery and were not moved again for at least one week. The tetrodes were advanced slowly (μm/day) towards the recording target, reaching either CA1 or mPFC after 2.5–3 weeks. The proximity of CA1 tetrodes to the pyramidal cell layer was estimated by sharp wave / ripple profiles (Csicsvari et al., 1999); Tetrodes were lowered to the mPFC (3–4 mm) by turning the microdrive screws 10.6 – 14.2 times. To ensure recording stability, the tetrodes were not adjusted for >= 24 hours before behavior testing and recording.

Recording apparatus: Multiunit spike and LFPs were acquired using either a Cheetah 160 (N=1 rat from saline/muscimol study) or a Digital Lynx SX (Neuralynx). For both systems, an electrode interface board connected tetrodes to a headstage containing unity gain amplifiers to minimize cable motion artifact; the headstages were connected to the system amplifiers with thin-wire tethers. Unit activity was filtered between 600 and 9000 Hz and digitized at 32,000 Hz (or about 28,000 Hz for the Cheetah 160) prior to online spike detection. For each tetrode, amplitude thresholds were manually set on each wire to maximize the signal to noise ratio for spike detection. When the amplitude on a single wire rose above the threshold, the signal waveform around the threshold crossing was saved (along with a timestamp) for all four wires. The waveforms were sorted offline into single units. LFPs signals were sampled continuously at 2000 Hz with a band-pass filter (1– 512 Hz) and recorded with the active electrode referenced directly to a skull screw implanted above the cerebellum or to another tetrode implanted in the brain. In the latter case, the reference tetrode was also recorded with respect to a skull screw implanted above the cerebellum so that the signal from the active electrode with respect to the skull could be recovered offline by subtraction. The position of LEDs mounted on the headstage was recorded by an overhead video camera, digitized (30 Hz), converted to time-stamped XY coordinates, and stored for offline analysis by the Cheetah recording system.

General surgical procedures: Each animal was given its daily allotment of rat chow ~2 hours before surgery, and was then anesthetized in a Plexiglas chamber with 5% isoflurane delivered at a rate of 1 L/min. Once deeply anesthetized, the rat was given ketoprofen subcutaneously (3 mg/kg; Sigma) to minimize post-operative pain, its head was shaved and positioned in a Kopf stereotax. Isoflurane (1–3%) delivered via a nose cone maintained anesthesia during surgery. The scalp was cleaned with Povidone-Iodine (Dynarex) and anesthetized with 0.7 cc of lidocaine / epinephrine (0.5% / 1:200,000; Hospira, Inc.). Core body temperature was monitored and maintained using a rectal probe and heating pad (part number: ATC 1000; World Precision Instruments). Sterile normal saline (1 cc) was delivered subcutaneously every hour to maintain hydration, and an ophthalmic ointment (Puralube vet ointment; Dechra) was applied to the animal's eyes and reapplied hourly; otherwise the eyes were covered for the duration of surgery. Twenty min after keotprofen administration, the scalp was resected and the skull was cleaned using distilled water and a dilute solution of hydrogen peroxide. Five holes were drilled into the thicker parts of the skull for stainless steel bone screws (part number: 40-77-8; FHC) to stabilize the implants. The skull was cleaned again and lambda and bregma were set to the same DV level. Target sites were measured and marked, thin layers of Metabond (Parkell) and Panavia (Kuraray) were applied to the skull to increase implant stability, excluding the areas around the implant site. Burr holes were drilled through the skull above the implant target site to expose the dura, and a stereo microscope was used to ensure that it was fully exposed and clean. The dura was incised with micro scissors, retracted, and the exposed cortex was kept clean and hydrated with normal saline until the implant was lowered into place. The craniotomy around the implant was sealed with Kwik-Cast (WPI), and the entire implant was then secured to the skull with dental acrylic (Coltene/Whaledent Inc.). After surgery, animals were returned to a clean home cage containing a wet mash of rat chow. For the three days following the procedure the mash was infused with a Meloxicam suspension (1 mg/kg; Boehringer Ingelheim) to minimize pain due to swelling, and the rats were monitored to ensure they consumed the food. Animals generally responded well to this treatment and return to preoperative levels of activity after three days. Animals were not handled by the experimenter for 7 days post-surgery.

Cannula implants: Cannula pre-cut to size were purchased from Plastics One. Guise cannula (26 gauge; C315GRL/SPC, 3.5 mm from plastic base to the tip), and the dummy and infusion cannula (33 gauge; C315DC, and C315I, respectively) extended 1.5 mm beyond the guide. The relatively short guide cannula was used to minimize damage to the frontal white matter during implantation. The tip of the infusion cannula targeted the mPFC (mm from bregma: AP: +3.0, ML +/− 0.92, and −3.45 mm ventral to the cortical surface)

Histology: After the recording experiments the rats were anesthetized with pentobarbital, perfused transcardially with normal saline followed by formalin, the brains were removed and cryoprotected in a PBS sucrose solution. Coronal slices (40um) cut by cryostat were taken to span the electrode target regions (AP from Bregma, mPFC = 2.5 – 4.0; CA1 = −3.0 – −5.5). The slices were stained with formol-thionin and the electrode tracts were inspected microscopically. The mPFC electrode bundles extended through infralimbic and prelimbic cortices with the final depth of the electrode tips above. The CA1 electrode bundles terminated above the stratum radiatum.

Replication: All analyses were carried out at the group level, and we have not replicated the tests in another group of animals.

Strategy for randomization and/or stratification: All experiments were performed sequentially using animals selected randomly from an available population; experiment assignments were not randomized. All behavioral experiments utilized a within-subjects design, and the order of conditions was fixed across animals.

Blinding: The experimenter was not blinded with respect to any task parameters or conditions.

Sample-size estimation: Our aim was to power each experiment in order to detect statistical effects with as few animals as possible. However, calculation of the required N's is difficult when true effect sizes are unknown, as was the case with these experiments. We selected N's based on our previous experience using within-subjects behavioral (e.g., Rich and Shapiro, 2009; Young and Shapiro, 2009)) and recording studies (e.g., (Riceberg and Shapiro, 2017)).

Inclusion and exclusion criteria: All animals that were trained to criteria in the behavioral task were included in their respective experiments.

QUANTIFICATION AND STATISTICAL ANALYSIS

Behavior

General behavioral analyses: Spatial task behavioral performance is illustrated in Figure S1. Maze behavior was categorized in four possible journey types: north-east (NE), north-west (NW), south-east (SE), and south-west (SW). During recording experiments, the rats' horizontal and vertical coordinates from the video tracker was linearized and expressed in terms of distance from the start of each journey. Reliable trajectories for all trials of a given journey type were fit with a polynomial ranging between 5th and 7th order to derive a canonical trajectory onto which data from individual trials were projected. Unit activity was analyzed on the start arm from 20 to 58 cm marks, after the rat turned toward the maze center and before it entered the choice point. Differences in heading angle and running speed did not influence the decoding or LFP analyses (see supplemental information Figure S2.2). Running speed was assessed on a trial-by-trial basis to generate a distribution of velocities at each point on the maze. Positions on the maze within a trial in which animal's speed fell outside of two standard deviations of the mean for that position were removal from LFP analyses.

Behavioral measures: Behavioral performance was quantified by three metrics: trials to criterion (TTC), number of errors (NE), and percent correct (PC). TTC was the number of trials the animal performed during the initial discrimination (ID) or any of the reversals (R1 or R2) prior to reaching the learning criterion (10 correct trials within a block of 12 consecutive trials). NE was the number of errors committed during any of these epochs, and PC was the proportion of correct trials during any of these epochs. Analyses carried out on TTC included only learning epochs in which animals reached the learning criterion, whereas NE and PC include any reversal block in which animals completed at least 12 trials.

Quantifying learning curves: We used the Smith algorithm (Smith et al., 2004) to derive learning curves (e.g., Figure S1) and to divide trials into early learning and stable performance epochs within reversals (Rich and Shapiro, 2009; Young and Shapiro, 2011). The Smith algorithm utilizes all trial outcomes across a learning epoch to calculate the probability (with confidence intervals) that a rat will select the correct goal on each trial. We defined early learning as the trials before the Smith algorithm predicted with 95% confidence that the animal was performing above chance. Stable performance was defined as the trials after this point. These operational definitions categorized trials for physiological analyses to compare interactions between the mPFC and the hippocampus during standard levels of learning and stable performance.

Recording Analyses

Movement analysis: The decoding analyses described in the main text showed that population codes in both mPFC and CA1 differentiated animals' pending goals on the start arm of the maze. To determine the extent to which the decoding could be explained by different behaviors on the start arm, we analyzed differences in heading angles in 1 cm increments along the segment of the start arm used for the decoding analysis (20 to 58 cm). The mean heading angle was calculated at each position for each journey type, and the differences between East and West journeys were calculated for the North and South start arms separately. A null difference distribution was generated by shuffling trial labels for each start arm and re-calculating the heading angle differences 1000 times. Visual inspection of the data revealed that the differences were von Mises distributed, and the parameters of the distribution were calculated accordingly (Fisher, 1993). The probability that the actual heading angle difference was obtained by chance alone was evaluated according to the null difference PDF, and the threshold for a significant difference was set at alpha = 0.05, FDR corrected at each position (see Figure S2.2).

Spike sorting: Spikes recorded on individual tetrodes were clustered into functional units, with single neurons as putative sources. Spike waveform parameters describing the shapes and relative amplitudes of the waveforms across all four tetrode wires were selected to define the basis for a space in which clustering was performed. Semi-automatic clustering was performed in several steps starting with KlustaKwik (http://klustakwik.sourceforge.net/). Noise clusters (e.g. from chewing artifact) were manually rejected, the remaining data automatic clustered again and then edited to identify well-segregated clusters of spikes that were assigned to functional units. The stability of individual units was assessed using a combination of waveform features and firing rates. The mean spike amplitude for a given unit was calculated in 20 equal temporal bins spanning the entire recording session for each tetrode wire. Any unit showing a significant Spearman rank correlation between time and amplitude on any wire ($\alpha = 0.05$, uncorrected) and an accompanying drift in mean firing rate ($\alpha = 0.05$, uncorrected) were rejected from further analysis. Visual inspection revealed that the statistical approach was conservative and units that seemed stable by eye were rejected; the included units had markedly stable waveforms.

Units with stable waveforms were clustered into putative pyramidal cell and interneuron groups separately for each region analyzed (i.e., CA1 and mPFC) based on waveform features (Bartho et al., 2004) and firing rate. Spike asymmetry and firing rate were the best discriminators, with interneurons having higher firing rates and more asymmetric waveforms and making up approximately 8% of the total number of identified units. All further analyses describe putative pyramidal units.

All analyses of neural representations (i.e., multidimensional scaling, SVM, and Granger) excluded units with firing rates < 2% of the unit with the highest firing rate. Units firing ≤ 20 spikes on the start arm were excluded from the spike-LFP phase locking analysis. The number of units included in each analysis is described in the relevant sections of the main text. The mean number of stable units recorded in single sessions were: Dual site, CA1 N = 29.4 (s.d. = 13.0), mPFC N = 17.2 (s.d. = 6.5); CA1 recordings with mPFC saline infusions: N = 28.1 (s.d. = 11.7); CA1 recordings with mPFC muscimol infusions: N = 27.2 (s.d. = 11.6).

Local field potentials

Pre-Processing: LFPs were segmented into single trial epochs for analysis and marked as useable or unusable on a trial-by-trial basis. CA1 LFPs were inspected for sharp waves to determine the relative placement of tetrodes with respect to the pyramidal cell layer (Csicsvari et al., 1999). For each recording session, one tetrode from mPFC and CA1 was selected for LFP analysis by the proportion of usable trials and the number of recorded units. For CA1, LFPs showing positive-going sharp waves were selected to ensure fair comparison of theta phase across recordings (Buzsaki, 2002). After re-referencing, the LFPs were filtered (2–250 Hz) for further analysis. To measure spike-phase locking LFPs were filtered for theta (5–18 Hz).

Frequency analysis: LFPs were padded with a time-reversed copy of themselves to prevent wrap-around effects due to convolution in the frequency domain and wavelet transformed using Morlet wavelets in 1–14 Hz, 14–30 Hz, and 30–80Hz frequency bands. The wavelet factor (i.e., the parameter controlling the tradeoff between frequency and time resolution) was set to 14 for 1–14 Hz, 24 for 14–30 Hz, and 32 for 30–80Hz. The position of the animal on the plus maze was interpolated at each time point, and time slices of the spectra were saved in 1 mm increments so that they could be averaged across trials.

Coherence: Position-locked, complex wavelet spectra were used to calculate coherence. Coherence measures the degree to which the phase and amplitude of one signal can be predicted by the same information from another. Perfect coherence ($C_{xy}2 = 1$; see below) occurs when the phase offset between two signals is constant across observations (here, across trials at a given position on a maze at a particular frequency) and the amplitude of both signals scale perfectly with one another. Zero coherence occurs when phase and amplitude vary randomly with respect to one another. Coherence was calculated as follows (Cohen, 2014):

$$C_{xy}{}^2 = \left| \frac{S_{xy}}{\sqrt{S_{xx} \times S_{yy}}} \right|^2$$

S_{xy} is the cross-spectral density and S_{xx} and S_{yy} are the autospectral densities. If the power spectrum of a given signal is normalized to unity, $C_{xy}2$ measures the consistency of the phase offset between the two signals. The manuscript reports raw coherence, and the supplemental information show that the coherence was driven almost entirely by phase consistency (Figure S3.1).

Spike-phase analysis: The preferred firing phase of units was measured with respect to the hippocampal theta oscillation by calculating the phase angle coincident with each spike from each neuron. The distribution of phase angles for each given unit was quantified by the Rayleigh test for non-uniformity. Filtered CA1 LFPs (5 – 18 Hz) were Hilbert transformed to extract the instantaneous phase of the theta oscillation, and the endpoints of each theta cycle were detected as phase ‘wraparounds’ from pi to −pi. The phase corresponding to each unit’s spike time was interpolated linearly between these endpoints to avoid the nonuniform distribution of Hilbert transform values caused by the saw tooth shape of theta oscillations.

Ensemble Coding and the “Kernel trick”: We used nonlinear rate vectors to investigate the similarity of neural representations across trials as animals used different memories to guide behavior. The elements of these vectors include not only the firing rates of individual neurons, but also the products of firing rates of pairs of neurons, triplets, or higher order groupings (the method for selecting the size of the groupings, referred to as the expansion order, is described in the methods section of the main text). A practical drawback to this method is that the number of elements in these vectors can become prohibitively large for explicit calculation even for moderate ensemble sizes and relatively low order expansions. The kernel method, colloquially referred to as ‘the kernel trick’ (Schölkopf, 2001; Schölkopf, 2002), provides a method via which the similarity amongst these expanded rate vectors can be calculated without having to explicitly compute the interaction terms of the vector for each trial. We used the inhomoengous polynomial kernel to calculate the similarity amongst expanded rate vectors for the analyses described in the main text. This is a general form of the simpler homogenous polynomial kernel function, described below.

The polynomial kernel takes the form:

$$k(\boldsymbol{x}, \boldsymbol{y}) = \langle \Phi(\boldsymbol{x}), \Phi(\boldsymbol{y}) \rangle$$

where $\langle .,. \rangle$ indicates the dot product, Φ represents the mapping of our simple rate vector into a space that includes products amongst individual units, $\boldsymbol{x}$ and $\boldsymbol{y}$ are our rate vectors for trials x and y, and $k(.,.)$ is the kernel (i.e., the similarity of population activity between the two trials). For this example, assume that ***x*** and ***y*** are two elements long (i.e., they include rates from just two units) and a second order expansion (i.e., one that includes only pairwise products of units). The similarity amongst these trials including the expanded rate vectors is obtained by calculating the dot product between ***x*** and ***y*** raised to the power of 2, the expansion order in this example.

$$k\left(\boldsymbol{x},\boldsymbol{y}\right)=\langle\boldsymbol{x},\boldsymbol{y}\rangle^{2}=(x_1y_1+x_2y_2)^{2}=x_1^2y_1^2+x_2^2y_2^2+2x_1x_2y_1y_2$$

The expanded rate vector for ***x*** (which need not be explicitly calculated) would be:

$$\Phi\left(\boldsymbol{x}\right)=\left(x_1^2,x_2^2,\ \sqrt{2}x_1x_2\right)$$

This expansion includes all pairwise products of firing rates.

The inhomogenous polynomial kernel takes the form:

$$k\left(\boldsymbol{x},\boldsymbol{y}\right)=\langle\Phi\left(\boldsymbol{x}\right),\Phi\left(\boldsymbol{y}\right)\rangle=(c+\langle\boldsymbol{x},\boldsymbol{y}\rangle)^{p}$$

The inhomogenous polynomaial kernel includes a constant term, *c*, that controls the inclusion of polynomial terms *up to* order *p*, whereas the homogenous polynomial kernel ($c = 0$) includes *only* polynomial terms of order *p*. Setting $c = 1$ for the example above, the calculation becomes:

$$k\left(\boldsymbol{x},\boldsymbol{y}\right)=(c+\langle\boldsymbol{x},\boldsymbol{y}\rangle)^{p}=1+2x_1y_1+2x_2y_2+2x_1x_2y_1y_2+x_1^2y_1^2+x_2^2y_2^2$$

The equivalent expanded rate vector would be:

$$\Phi\left(\boldsymbol{x}\right)=\left(1,\ \sqrt{2}x_1,\ \sqrt{2}x_2,\ \sqrt{2}x_1x_2,x_1^2,x_2^2\right)$$

This representation includes terms related to the individual unit firing rates ($\sqrt{2}x_1,\ \sqrt{2}x_2$), as well as the pairwise products of firing rates ($\sqrt{2}x_1x_2,x_1^2,x_2^2$). As the expansion order and the number of units in the simple rate vector increase, the number of elements in the expanded rate vector becomes very large, thus the appeal of the kernel trick. We set $c = 1$ for all of our analyses so that expansions included polynomial terms up to and including the expansion order. For example, if the expansion order was set to 3, then the expanded rate vector includes the individual unit rates, all pairwise products of individual units, and all triplet-wise products.

Ensemble activity was normalized before further processing. Z-scoring the firing rates of individual units prevents units with high activity from having disproportionate influence on classification, and normalization of rate vectors limits the maximum similarity to one value so that no trial is more self-similar than any other. Expansions up to 8th order were simultaneously optimized with respect to KKT complementary condition violation tolerance, box constraint, and kernel scale (see below) using a direct search procedure. The *box constraint* parameter affects the 'softness' of the SVM margin by controlling the penalty for misclassifications and is used to minimize over-fitting. The *kernel scale* parameter divides the kernel arguments by a scalar so that values in the kernel matrix don't become so large or so small that round off errors occur due to limitations on machine precision. The *KKT complementary condition violation tolerance* value sets the stopping tolerance for the algorithm. The direct search procedure explores the parameter space by iteratively moving along directions that best minimize a cost function, which in this case was the cross-validation misclassification rate. The direct search algorithm was run with random start points 100 times for each expansion order, and the parameter set with the best cross validation accuracy and lowest order expansion (in the case of a tie) was chosen for subsequent analysis.

Neuronal trajectories—To quantify population coding during learning and memory performance, we calculated the similarity of expanded population vectors for all pairs of trials in a 38 cm section of each start arm where rats approached but did not reach the choice point (20–58 cm of the linearized journey). Firing rates of each unit were calculated for each trial in the start arm, z-score normalized across trials, and smoothed with a Gaussian of ½ trial standard deviation. Smoothing was carried out on a per-neuron basis across trials by first padding the trial-wise rate vector with a time-reversed copy of itself. Data from missing trials due to poor video tracking were linearly interpolated. The resulting padded rate vector was convolved with the Gaussian by multiplication in the frequency domain, and then transformed back to the time domain. The padded portion and interpolated elements were then discarded. The purpose of padding the rate vector with a time-reversed copy of itself was to prevent mixing of information from the first and last trials of the recording session.

The per-trial rate vectors included all simultaneously recorded units that met inclusion criteria and were normalized to unit length. To measure ensemble states during learning, all trials within the initial discrimination and each reversal were divided into 7 bins spanning the learning curve. Firing rate vectors recorded in trials within a bin were averaged, the similarity across all pairs of bins was computed using the kernel trick, and the unique pairs of the resulting dot product bins were reduced to 3 dimensions using multidimensional scaling. A 2-D projection of the 3-D embedding is shown in Figs. 3A and B.

Support vector machines—Support vector machines (SVMs) quantified the extent to which ensemble activity in mPFC and CA1 distinguished between animals' pending goal choices (Schölkopf, 2002). SVMs were trained to distinguish two classes of data ('go east' vs. 'go west') using Matlab's statistics toolbox (R2014a, Mathworks). For each ensemble, single unit firing rates were z-score normalized across trials, and each trial's vector was normalized to unit length. SVMs were fit using an inhomogeneous polynomial kernel, and each SVM was optimized to maximize the leave-one-out cross validation accuracy (see below and supplemental information). The mean expansion order required for CA1 goal prediction in the dual-site implant animals was 2.64 (s.d. = 1.75), and 3.00 (s.d. = 1.92) for mPFC. The average expansion order for CA1 goal prediction after mPFC saline infusion was 2.90 (s.d. = 2.45), and after muscimol was 2.75 (s.d. = 2.42). The average expansion order required for CA1 start arm decoding after mPFC saline infusions was 1.05 (s.d. = 0.22) and after saline infusions was 1.00 (s.d. = 0.00).

Cross validation accuracy and significance testing—The accuracy of each SVM was defined as the leave-one-out cross validation accuracy. A given SVM was fit to a dataset with the exception of a single trial, which was subsequently input to the SVM to classify. The cross validation accuracy was the proportion of trials correctly classified. To determine if the SVMs were classifying trial goals by discovering task structure or unrelated noise, we repeated the above procedure using the same data and parameter sets after shuffling the trial labels (e.g., the current goal or current start arm). We calculated the cross-validation decoding accuracy following 1000 separate shuffles of the data to determine if the observed decoding accuracy was better than chance, defined here as > 95% of the shuffles.

Control procedure: SVM decoding of time-shifted goal vectors: Previous reports note that stably recorded units exhibit firing rates that drift over time (Hyman et al., 2012), which we also observed in our trajectory plots (Figure 3). To determine if the SVMs were discovering neural states associated with temporal variation in firing rates as a proxy for task features, we circularly shifted trial label vectors in time (e.g. the last trial became the first, and so on), until we found a shifted version that exhibited a near 0 rank correlation with the original (+/− 0.100) and re-evaluated the SVM using the shifted vector. Unlike the shuffled cross validation, the shifted vector has similar temporal properties as the original. We found that decoding of the time shifted data was almost always worse than those with the actual goal or start arm labels, indicating that the SVMs were not discovering convenient temporal structure that served as a proxy for current goals or start arms. For the dual implant animals, decoding accuracy for all 28 CA1 ensembles was greater using the original goal state vector (expected: 14 out 28; χ^2 (1) =14.0, $p < 0.05$), and 19 out of 28 of the SVMs fit to time shifted data showed better than chance decoding. For mPFC, 27 out of 28 ensembles had greater decoding accuracy of the actual goal-state relative to the time-shifted version (expected: 14 out of 28; χ^2 (1) = 12.07, $p < 0.05$), and 18 out of 28 of the SVMs fit to time shifted data showed better than chance decoding. For CA1 recordings along with ipsilateral mPFC saline infusions, 20 out of 21 ensembles had better decoding accuracy on the original goal vector (expected: 10.5 out of 21; χ^2 (1) = 8.60, $p < 0.05$), and every ensemble was better for the muscimol condition (12 out of 12; expected: 6 out of 12; χ^2 (1) = 6.00, $p < 0.05$). Similarly, decoding of the current start arm was better for data with actual start arm labels relative to their time-shifted counterparts (saline: 21 out of 21; expected: 10.5 out of 21; χ^2 (1) = 10.50, $p < 0.05$; muscimol: 12 out of 12, expected: 6 out of 12; χ^2 (1) = 6.00, $p < 0.05$).

SVM margins—The SVM margin width quantified the magnitude of goal discrimination by ensembles using the SVM score, a signed value indicating the distance of a trial from the decision boundary. The SVM margin is the difference in SVM scores between a given trial (the ‘true’ class) and the trial of the opposite class (the ‘false’ class) with the largest SVM score and thereby estimates the distance or separation between two classes. Because each SVM was fit using parameters optimized for a given ensemble, we normalized the margins before subsequent analysis. A null distribution of SVM margins was estimated from models fit using shuffled trial labels, and actual SVM margins were z-scored according to the null distribution. Because the normalized margins were not normally distributed, the data were log transformed before ANOVAs or t-tests to stabilize the variance. Only trials in which animals made correct responses were included in analyses of SVM margins.

Granger prediction—Granger prediction assesses the temporally directed statistical relationship between two time series by testing if the recent history of one time series predicts changes in a target time series beyond what is predicted by the history of the target series itself. The *Granger value* quantifies this prediction as the log ratio of the residual variances for the model incorporating only one time series (the target) and the model incorporating both. The higher the Granger value, the greater the prediction gained by including the second time series. Firing rates of individual units were z-score normalized across trials and smoothed (1/2 trial standard deviation) to minimize variance due to place coding. The smoothing procedure involved linearly interpolating missing trials (e.g., due to poor video tracking), padding individual units' trial-wise rate vectors with a time-reversed copy of itself, and convolution with a Gaussian via multiplication in the frequency domain. The result was then transformed back into the time domain, the padded portion of the vector was discarded, and interpolated data points removed. The padding procedure prevented mixing of information from the first and last trials of the recording session. The activity state of each trial was defined by the dot product of its population vector with the mean vector of the last two trials of the previous contingency using the kernel trick. Dependent variables in the linear models were power transformed to stabilize the variance.

Granger predictions were compared to chance using permutation testing. After calculating the residual variance of the model that incorporated only the target time series (the numerator of the Granger value), we shuffled the trial order of the non-target series and calculated the Granger value after each of 1000 shuffles to generate a null distribution. Shuffling eliminated the temporal correspondence between the two series and tested the extent to which any increase in explained variance of the target was produced merely by the increased number of independent variables provided by the non-target series. Actual Granger values were compared to the shuffled values to calculate the corresponding probability. A similar procedure was followed to assess differences between Granger values. In this procedure, the within-trial correspondence between variables in an observation (e.g. $CA1_{t-1}$, $CA1_{t-2}$, $mPFC_{t-1}$, $mPFC_{t-2}$) was preserved, and observations were shuffled between models (e.g. from the model examining the influence of mPFC on CA1 to the model examining the influence of CA1 on mPFC) prior to calculating Granger values. The difference in Granger values was then calculated after each shuffle to generate the null difference distribution against which the observed difference was compared.

Granger prediction: Bayes information criterion: The dependent variable (DV) for each observation was the similarity of a given trial's ensemble activity to that observed at the end of the previous learning epoch (i.e., ID, R1, or R2). The independent variables (IVs) were the similarity measures for the two previous trials for either the target time series alone or both time series. The number of IVs (2) was determined by minimizing the Bayes' information criterion (BIC) with respect to trial number. The BIC aids in model selection by balancing model fit with parsimony (i.e., the number of parameters). The BIC was calculated using the following equation:

$$BIC = n * \ln(RSS/n) + k * \ln(n)$$

where n is the number of observations in the model, RSS is the residual sum of squares, and k is the number of parameters in the model that includes both time series.

Granger prediction: expansion order: Granger prediction was carried out on neuronal ensemble states calculated using expanded rate vectors, as described for the SVM analysis. A third order expansion was selected for every ensemble to ensure that all data were represented on the same scale. This order expansion was chosen because it was the average order expansion required to maximize decoding accuracy in the SVM analysis.

Statistical analyses: All statistical analyses of behavioral data were carried out using standard ANOVAs and t-tests. All correlations were carried out using the Pearson correlation method unless otherwise noted (e.g., Spearman rank correlation for evaluating randomized task parameters). Where parametric statistical tests were not appropriate (e.g. in comparing Granger values), permutation techniques were used. Statistical significance was set at alpha = 0.05 for all statistical analyses unless otherwise noted.

DATA AND SOFTWARE AVAILABILITY

Data and MATLAB scripts are available upon request from Matthew Shapiro (matthew.shapiro@mssm.edu)

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Chemicals, Peptides, and Recombinant Proteins		
Forane (Isoflurane)	Baxter	NDC: 10019-360-60
Ketoprofen	Sigma	K1751
Lidocaine / Epinepherine	Hospira	NDC: 0409-0996-01
Puralube Vet Ointment	Dechra	NDC: 17033-211-38
Metabond	Parkell	S380
Panavia	Kuraray	488KA
Muscimol	Sigma	M1523
Experimental Models: Organisms/Strains		
Long Evans Rats	Charles River	006
Software and Algorithms		
MATLAB R2014a	Mathworks	R2014a
Other		
96 Channel DigitalLynx 16SX	Neuralynx	neuralynx.com/products/digital_data_acquisition_systems/digital_lynx_16sx

REAGENT or RESOURCE	**SOURCE**	**IDENTIFIER**
Cheetah 160	Neuralynx	neuralynx.com
12.5 μm nichrome wire	Kanthal Precision Technology	RO-800
NI-DAQ USB-6008	National Instruments	779051-01
Stainless steel bone screws	FHC	40-77-8
Guide cannula	Plastics One	C315GRL/SPC
Internal cannula	Plastics One	C315I
Dummy cannula	Plastics One	C315DC
Metabond	Parkell	S380
Panavia	Kuraray	488KA

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank Justin Riceberg, PhD, Pierre Enel, PhD, Pablo Martin, Vincent Luo, and Phil Avigan for helpful comments on the manuscript, and Jordana Goldman, Anna Balk, and Maojuan Zhang, MD for technical help. This material is based on work supported by NIMH grants 5R01MH065658, 5R01MH073689, and 5P50MH094263, the National Science Foundation Graduate Research Fellowship under Grant No. 1443116, the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program, the Icahn School of Medicine at Mount Sinai, and the NVIDIA corporation.

References

Bartho P, Hirase H, Monconduit L, Zugaro M, Harris KD, Buzsaki G. Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. Journal of neurophysiology. 2004; 92:600–608. [PubMed: 15056678]

Bartlett, FC. Remembering: A study in experimental and social psychology. Cambridge, UK: Cambridge University Press; 1932.

Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. Journal of the Royal Statistical Society, Series B. 1995; 57:289–300.

Bergström. Experiments upon Physiological Memory by Means of the Interference of Associations. The American Journal of Psychology. 1893; 5:356–369.

Braver TS. The variable nature of cognitive control: a dual mechanisms framework. Trends in cognitive sciences. 2012; 16:106–113. [PubMed: 22245618]

Buzsaki G. Theta oscillations in the hippocampus. Neuron. 2002; 33:325–340. [PubMed: 11832222]

Cohen, MX. Analyzing Neural Time Series. Cambridge, MA: MIT Press; 2014.

Cassaday HJ, Nelson AJ, Pezze MA. From attention to memory along the dorsal-ventral axis of the medial prefrontal cortex: some methodological considerations. Frontiers in systems neuroscience. 2014; 8:160. [PubMed: 25249948]

Cohen, MX. Analyzing Neural Time Series. Cambridge, MA: MIT Press; 2014.

Csicsvari J, Hirase H, Czurko A, Mamiya A, Buzsaki G. Fast network oscillations in the hippocampal CA1 region of the behaving rat. The Journal of neuroscience : the official journal of the Society for Neuroscience. 1999; 19:RC20. [PubMed: 10436076]

Depue BE. A neuroanatomical model of prefrontal inhibitory modulation of memory retrieval. Neuroscience and biobehavioral reviews. 2012; 36:1382–1399. [PubMed: 22374224]

Ferbinteanu J, Shapiro ML. Prospective and retrospective memory coding in the hippocampus. Neuron. 2003; 40:1227–1239. [PubMed: 14687555]

Fisher, NI. Statistical Analysis of Circular Data. New York, NY: Cambridge University Press; 1993.

Fujisawa S, Buzsaki G. A 4 Hz oscillation adaptively synchronizes prefrontal, VTA, and hippocampal activities. Neuron. 2011; 72:153–165. [PubMed: 21982376]

Goto Y, Grace AA. Dopamine modulation of hippocampal-prefrontal cortical interaction drives memory-guided behavior. Cerebral cortex. 2008; 18:1407–1414. [PubMed: 17934187]

Greenberg R, Underwood BJ. Retention as a function of stage of practice. Journal of experimental psychology. 1950; 40:452–457. [PubMed: 15436941]

Harlow HF. The formation of learning sets. Psychological review. 1949; 56:51–65. [PubMed: 18124807]

Hoover WB, Vertes RP. Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. Brain structure & function. 2007; 212:149–179. [PubMed: 17717690]

Hyman JM, Ma L, Balaguer-Ballester E, Durstewitz D, Seamans JK. Contextual encoding by ensembles of medial prefrontal cortex neurons. Proceedings of the National Academy of Sciences of the United States of America. 2012; 109:5086–5091. [PubMed: 22421138]

Ito HT, Zhang SJ, Witter MP, Moser EI, Moser MB. A prefrontal-thalamo-hippocampal circuit for goal-directed spatial navigation. Nature. 2015; 522:50–55. [PubMed: 26017312]

Jadhav SP, Rothschild G, Roumis DK, Frank LM. Coordinated Excitation and Inhibition of Prefrontal Ensembles during Awake Hippocampal Sharp-Wave Ripple Events. Neuron. 2016; 90:113–127. [PubMed: 26971950]

Johnson A, Redish AD. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2007; 27:12176–12189. [PubMed: 17989284]

Jonides J, Nee DE. Brain mechanisms of proactive interference in working memory. Neuroscience. 2006; 139:181–193. [PubMed: 16337090]

Kajikawa Y, Schroeder CE. How local is the local field potential? Neuron. 2011; 72:847–858. [PubMed: 22153379]

Kyd RJ, Bilkey DK. Prefrontal cortex lesions modify the spatial properties of hippocampal place cells. Cerebral cortex. 2003; 13:444–451. [PubMed: 12679291]

Lansink CS, Meijer GT, Lankelma JV, Vinck MA, Jackson JC, Pennartz CM. Reward Expectancy Strengthens CA1 Theta and Beta Band Synchronization and Hippocampal-Ventral Striatal Coupling. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2016; 36:10598–10610. [PubMed: 27733611]

McFarland WL, Teitelbaum H, Hedges EK. Relationship between hippocampal theta activity and running speed in the rat. Journal of comparative and physiological psychology. 1975; 88:324–328. [PubMed: 1120805]

McKenna JT, Vertes RP. Afferent projections to nucleus reuniens of the thalamus. The Journal of comparative neurology. 2004; 480:115–142. [PubMed: 15514932]

Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. Annual review of neuroscience. 2001; 24:167–202.

Navawongse R, Eichenbaum H. Distinct pathways for rule-based retrieval and spatial mapping of memory representations in hippocampal neurons. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2013; 33:1002–1013. [PubMed: 23325238]

O'Keefe JDJ. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain research. 1971; 34:171–175. [PubMed: 5124915]

Peters GJ, David CN, Marcus MD, Smith DM. The medial prefrontal cortex is critical for memory retrieval and resolving interference. Learning & memory. 2013; 20:201–209. [PubMed: 23512936]

Preston AR, Eichenbaum H. Interplay of hippocampus and prefrontal cortex in memory. Current biology : CB. 2013; 23:R764–773. [PubMed: 24028960]

Preuss TM. Do rats have prefrontal cortex? The rose-woolsey-akert program reconsidered. Journal of cognitive neuroscience. 1995; 7:1–24. [PubMed: 23961750]

Rainer G, Rao SC, Miller EK. Prospective coding for objects in primate prefrontal cortex. The Journal of neuroscience : the official journal of the Society for Neuroscience. 1999; 19:5493–5505. [PubMed: 10377358]

Rempel-Clower NL, Zola SM, Squire LR, Amaral DG. Three cases of enduring memory impairment after bilateral damage limited to the hippocampal formation. The Journal of neuroscience : the official journal of the Society for Neuroscience. 1996; 16:5233–5255. [PubMed: 8756452]

Riceberg JS, Shapiro ML. Reward stability determines the contribution of orbitofrontal cortex to adaptive behavior. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2012; 32:16402–16409. [PubMed: 23152622]

Riceberg JS, Shapiro ML. Orbitofrontal Cortex Signals Expected Outcomes with Predictive Codes When Stable Contingencies Promote the Integration of Reward History. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2017; 37:2010–2021. [PubMed: 28115481]

Rich EL, Shapiro M. Rat prefrontal cortical neurons selectively code strategy switches. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2009; 29:7208–7219. [PubMed: 19494143]

Rich EL, Shapiro ML. Prelimbic/infralimbic inactivation impairs memory for multiple task switches, but not flexible selection of familiar tasks. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2007; 27:4747–4755. [PubMed: 17460087]

Schölkopf, B. The kernel trick for distances. In: Leen, TKDTG., Tresp, V., editors. Neural Information Processing Systems. MIT Press; 2001. p. 301-307.

Schölkopf, BS., AJ. Learning with kernels: support vector machines, regularization, optimization, and beyond. Cambridge, MA: MIT Press; 2002.

Shannon CE. A Mathematical Theory of Communication. The Bell System Technical Journal. 1948; 27:379–423.

Shimamura AP, Jurica PJ, Mangels JA, Gershberg FB, Knight RT. Susceptibility to Memory Interference Effects following Frontal Lobe Damage: Findings from Tests of Paired-Associate Learning. Journal of cognitive neuroscience. 1995; 7:144–152. [PubMed: 23961820]

Sirota A, Montgomery S, Fujisawa S, Isomura Y, Zugaro M, Buzsaki G. Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. Neuron. 2008; 60:683–697. [PubMed: 19038224]

Smith AC, Frank LM, Wirth S, Yanike M, Hu D, Kubota Y, Graybiel AM, Suzuki WA, Brown EN. Dynamic analysis of learning in behavioral experiments. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2004; 24:447–461. [PubMed: 14724243]

Spellman T, Rigotti M, Ahmari SE, Fusi S, Gogos JA, Gordon JA. Hippocampal-prefrontal input supports spatial encoding in working memory. Nature. 2015; 522:309–314. [PubMed: 26053122]

Squire LR. Memory systems of the brain: a brief history and current perspective. Neurobiology of learning and memory. 2004; 82:171–177. [PubMed: 15464402]

Tulving E. Memory and consciousness. Canadian Psychology / Psychologie Canadienne. 1984; 26:1–12.

Underwood BJ. Interference and forgetting. Psychological review. 1957; 64:49–60. [PubMed: 13408394]

Vertes RP, Hoover WB, Do Valle AC, Sherman A, Rodriguez JJ. Efferent projections of reuniens and rhomboid nuclei of the thalamus in the rat. The Journal of comparative neurology. 2006; 499:768–796. [PubMed: 17048232]

Xu W, Sudhof TC. A neural circuit for memory specificity and generalization. Science. 2013; 339:1290–1295. [PubMed: 23493706]

Young JJ, Shapiro ML. Double dissociation and hierarchical organization of strategy switches and reversals in the rat PFC. Behavioral neuroscience. 2009; 123:1028–1035. [PubMed: 19824768]

Young JJ, Shapiro ML. Dynamic coding of goal-directed paths by orbital prefrontal cortex. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2011; 31:5989–6000. [PubMed: 21508224]

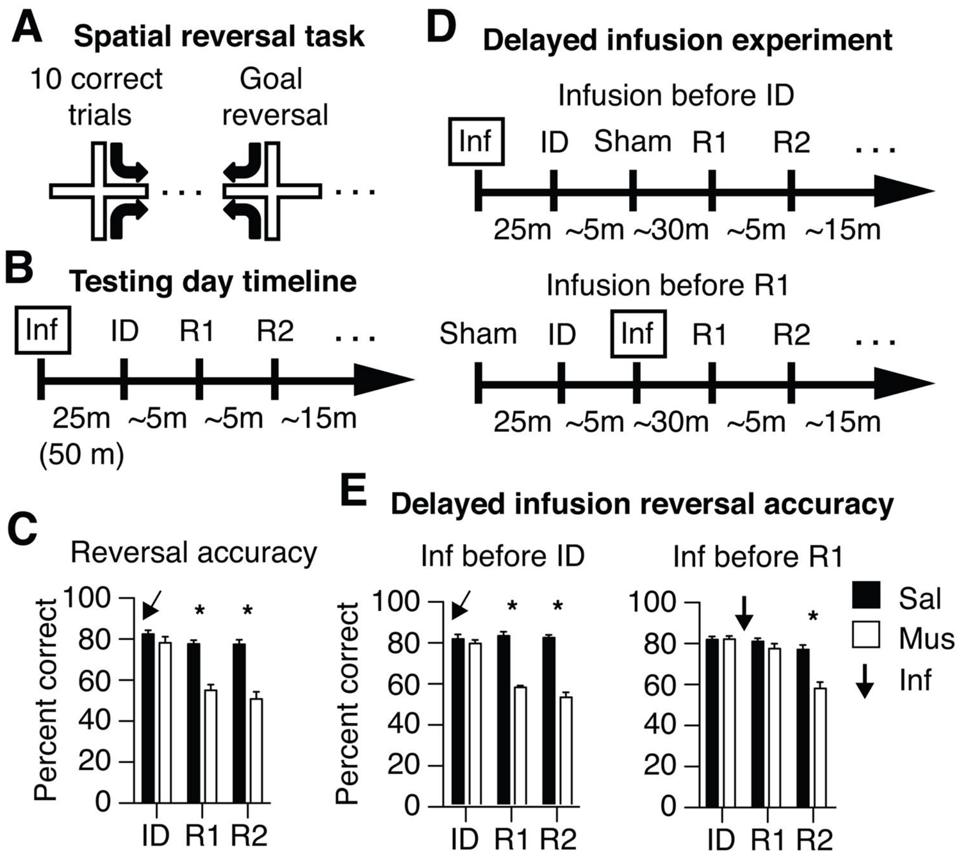

Figure 1.

mPFC inactivation impairs serial reversal learning in a hippocampus-dependent spatial memory task. (**A**) Rats are put on one of two pseudorandomly chosen start arms (North or South) in each trial and learn to associate one of two possible goal arms (East or West) with food reward. The same goal is rewarded during a block of trials until a rat performs well (10 of 12 correct choices), and then the other goal is rewarded in a new block of trials. Well-trained animals learned the initial discrimination (ID) and each spatial reversal in ~5 min. (**B**) The mPFC of each rat was infused bilaterally with muscimol or saline 25 or 50 minutes before testing. (**C**) mPFC inactivation had no effect on initial learning but profoundly impaired subsequent reversal learning (R1, R2). Similar effects were observed using other behavioral metrics, as illustrated in Figure S1 and Table S1A. (**D**) Additional experiments included a sham infusion and a 25 min delay between ID and R1. One experiment gave the real infusion before the ID and the sham infusion before R1 (upper timeline); a second experiment reversed this order so that the sham preceded the ID and the actual infusion preceded R1 (lower timeline). (**E**) Inactivating the mPFC before ID impaired R1 despite the additional delay (cf. panel C). Inactivating the mPFC after the ID and before R1 had no effect on R1, but impaired learning a second reversal (R2). Similar findings were obtained using other behavioral metrics, as illustrated in Tables S1B, S2A, and S2B. *above columns indicate significant differences between drug conditions within each learning epoch.

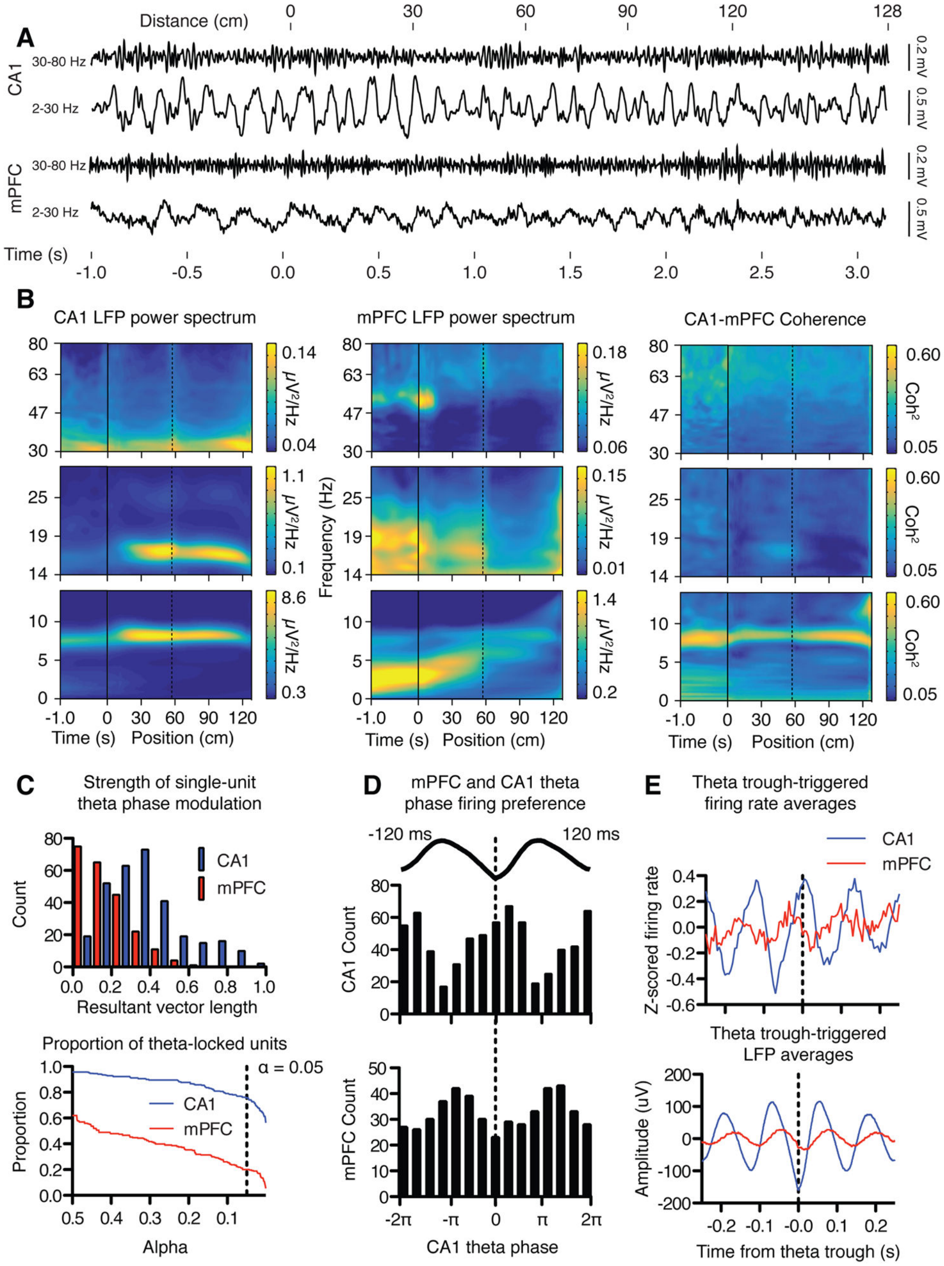

Figure 2.

Theta oscillations were coordinated in the mPFC and CA1 as rats perform the spatial reversal task. (**A**) Typical examples of LFPs recorded in CA1 (top traces) and mPFC (bottom traces) as rats progress through a single trial, each filtered in either low (lower trace) or high (upper trace) frequency bands. CA1 LFPs oscillate persistently in theta while mPFC LFPs revealed a different pattern in which oscillations increased in frequency and decreased in amplitude across the trial. (**B**) Power spectrograms quantified group averaged LFP data in 0–14, 14–30, and 30–80 Hz bands, and show that CA1 theta power was greatest on the start arm (left). Spectrograms verified the more dynamic mPFC power spectra (middle). Trials began with prominent 4 Hz oscillations that increased in frequency as animals approached the goal arm. Theta oscillations became prominent as animals entered the goal arm, and increased in frequency just before animals reached the reward point. Theta coherence was strongest before a trial was initiated, and after animals crossed into the goal arm (right). Negative values on the horizontal axis indicate time before initiation of the trial, and positive values indicate position on the maze in cm. The dotted line marks the boundary between the start arm and the maze choice point. Different panels represent wavelet transformations carried out using different parameters to optimize wavelet time-frequency resolution. (**C**) Unit activity in both structures is coordinated by hippocampal theta rhythm. A significant portion of units in both mPFC (20.2%) and CA1 (75.8%) showed greater than chance modulation of spiking activity by theta phase. Top: the magnitude of theta locking expressed in terms of the resultant vector length (see methods). Bottom: the proportion of CA1 and mPFC units passing significance at various alpha values. The dotted line represents alpha = 0.05. (**D**) mPFC and CA1 units fired at different preferred theta phases. Most CA1 units tended to fire near the trough of the hippocampal theta cycle, while mPFC units tended to prefer to fire early in the descending phase (Beta2 in (Lansink et al., 2016)). The different phase preferences correspond to a temporal offset in spiking between the two structures of about 51 ms. (**E**) Theta trough-triggered averages of the firing rate of mPFC and CA1 units verified the phase modulation of firing rates in both regions and the temporal offset of activity between the two structures.

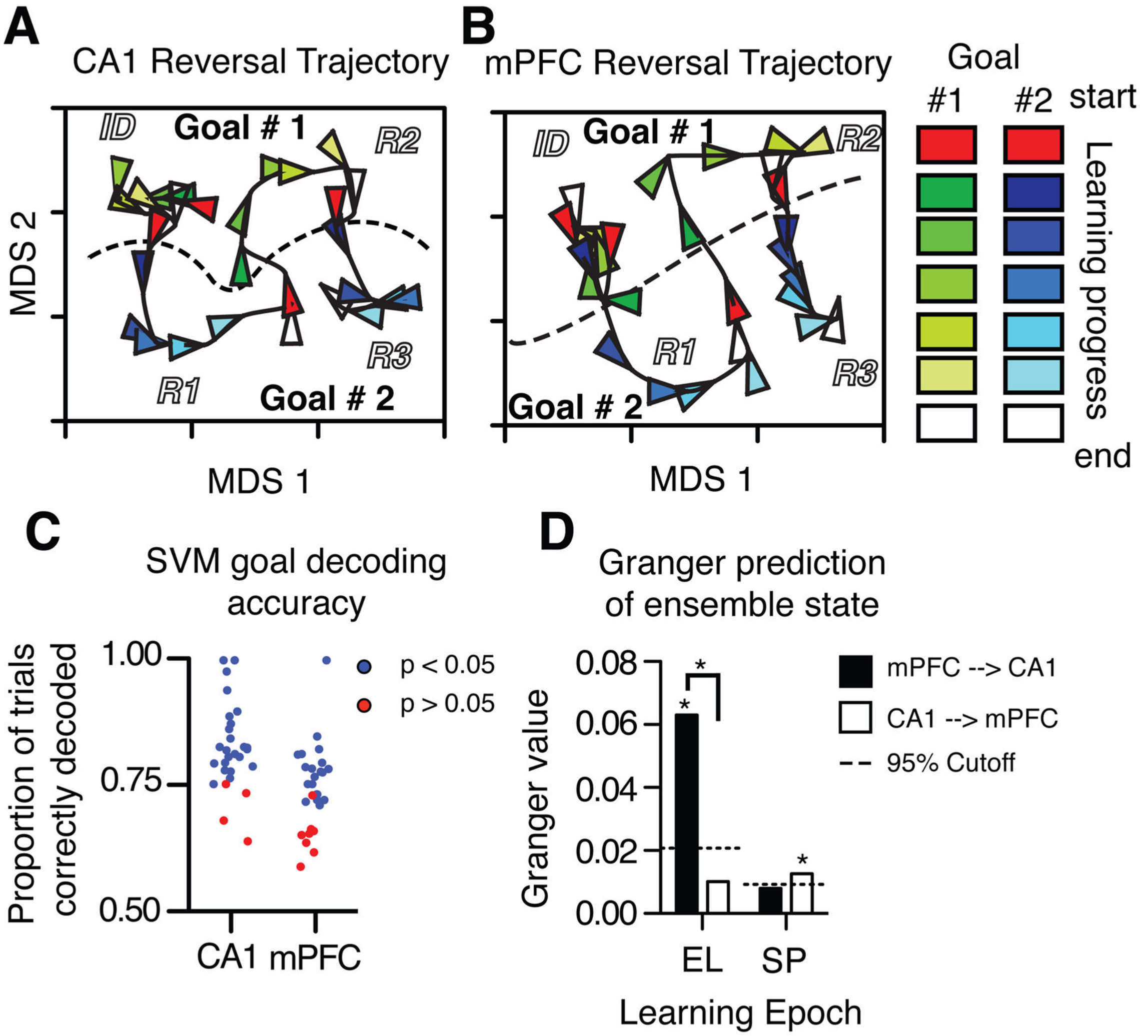

Figure 3.

Ensemble activity in both CA1 and mPFC distinguish rules, and mPFC activity predicts CA1 dynamics. All activity was recorded in the start arm before the rat entered the choice point. Population coding dynamics across testing sessions are illustrated qualitatively as 2D projections of CA1 (**A**) and mPFC (**B**) ensemble states analyzed by MDS. Trials in each contingency block were averaged and subdivided into 7 points to span learning curves, shown here by arrow heads. The colors show the progression of neuronal activity from the start of a new contingency (red) through the learning curve (brightening colors, goal 1 in shades of green, and goal 2 in shades of blue) to criterion performance and the end of the block (white). Ensemble states distinguish time (or distance) and goals (X and Y axes, respectively; goal 1 top, goal 2 bottom). (**C**) Ensemble codes recorded in the start arms were quantified by SVMs that decoded pending goals in single trials, showing rule coding in both brain regions. Each dot represents a single ensemble, blue indicates better than chance decoding. (**D**) The history of mPFC activity helped predict changes in CA1 activity early in learning (EL), but not during stable performance (SP), when CA1 activity has a small but significant influence on mPFC states. * Above columns indicate significantly different than chance defined by permutation tests.

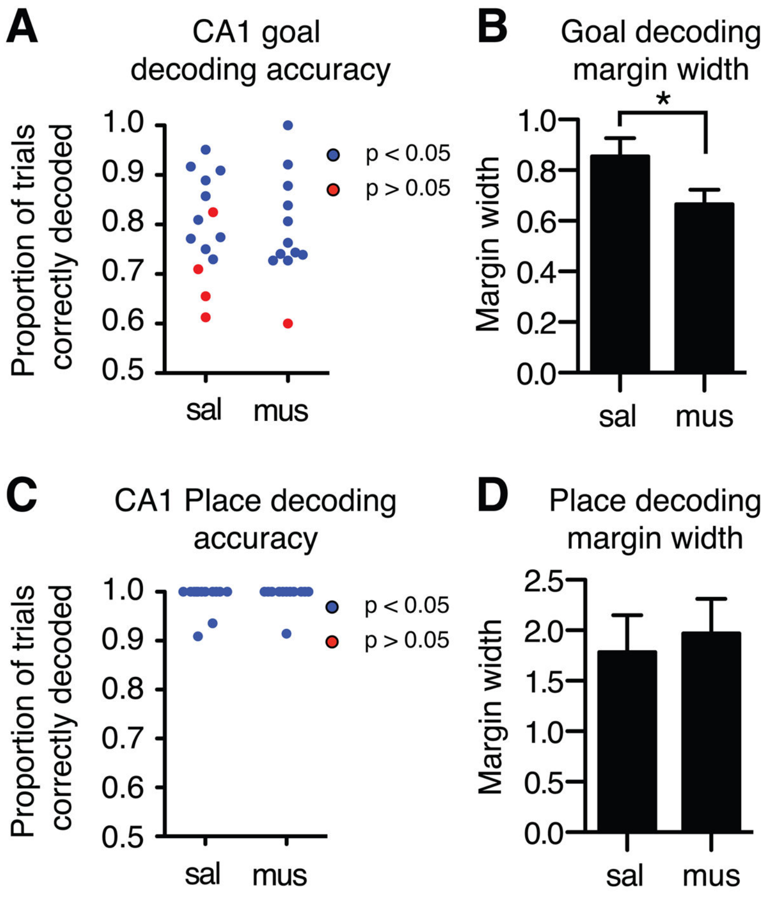

Figure 4.

mPFC inactivation reduces the separation of rule, but not place representations in CA1 ensembles recorded in the start arms. (**A**) mPFC inactivation did not impair prediction accuracy coding by CA1 ensembles, but (**B**) did reduced the separation between prospective codes as measured by the distance between activity patterns in each trial and the SVM goal-decoding margin. (**C**) mPFC inactivation did not affect start arm decoding accuracy by CA1 ensembles, (**D)** and did not alter the separation between spatial codes recorded in the start arm.

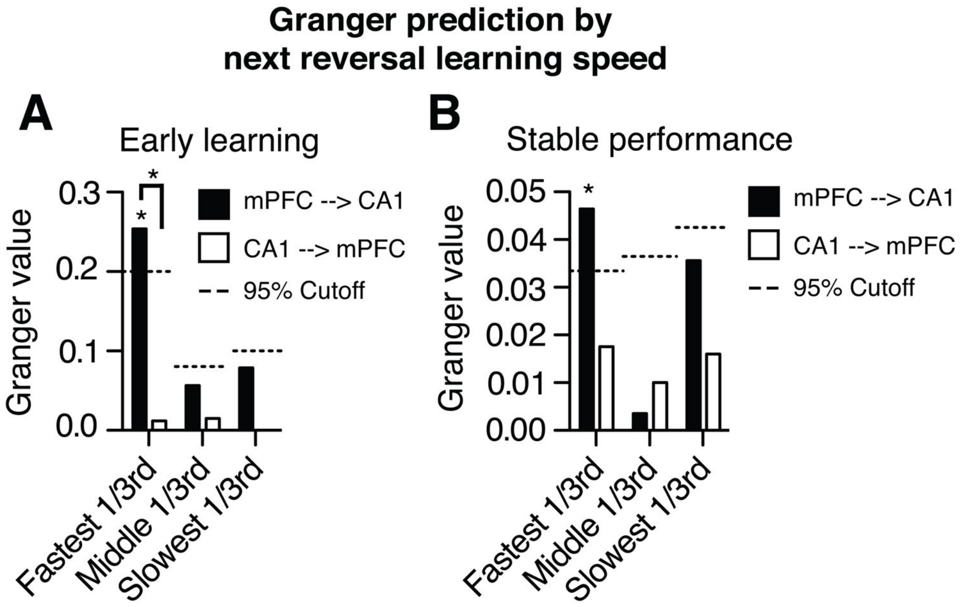

Figure 5.

mPFC activity during learning influenced CA1 activity most strongly when the subsequent reversal is learned most quickly. Granger values were calculated during early learning (**A**) and stable performance (**B**) for trials that preceded reversals that were learned at different rates. Significant influence of mPFC on CA1 was only observed when animals learned the subsequent reversal quickly, in ≤ 3 trials (see main text). * above columns indicate significantly different than chance calculated by permutation tests.