Article

Neuron

Replay of Behavioral Sequences in the Medial Prefrontal Cortex during Rule Switching

Highlights

- mPFC represents relative maze positions, forming a generalized code
- During immobility periods, the mPFC replays trajectories of generalized positions
- Trajectory replay across the mPFC and hippocampus is not coordinated
- mPFC trajectory replay positively correlates with rule-switching performance

Authors

Karola Kaefer, Michele Nardin, Karel Blahna, Jozsef Csicsvari

Correspondence

jozsef.csicsvari@ist.ac.at

In Brief

Kaefer et al. performed simultaneous mPFC and hippocampus recordings in rats solving a mPFC-dependent rule-switching task. They found that the mPFC reactivates spatial trajectories during awake immobility, which was independent of hippocampal replay and positively correlated with rule-switching performance.

Kaefer et al., 2020, Neuron *106*, 154–165
April 8, 2020
https://doi.org/10.1016/j.neuron.2020.01.015

CellPress

CellPress

Neuron
Article

Replay of Behavioral Sequences in the Medial Prefrontal Cortex during Rule Switching

Karola Kaefer,[1] Michele Nardin,[1] Karel Blahna,[1,2] and Jozsef Csicsvari[1,3,*]
[1]Institute of Science and Technology Austria (IST Austria), 3400 Klosterneuburg, Austria
[2]Present address: Laboratory of Experimental Neurophysiology, Biomedical Center, Faculty of Medicine in Pilsen, Charles University, 32300 Pilsen, Czech Republic
[3]Lead Contact
*Correspondence: jozsef.csicsvari@ist.ac.at
https://doi.org/10.1016/j.neuron.2020.01.015

SUMMARY

Temporally organized reactivation of experiences during awake immobility periods is thought to underlie cognitive processes like planning and evaluation. While replay of trajectories is well established for the hippocampus, it is unclear whether the medial prefrontal cortex (mPFC) can reactivate sequential behavioral experiences in the awake state to support task execution. We simultaneously recorded from hippocampal and mPFC principal neurons in rats performing a mPFC-dependent rule-switching task on a plus maze. We found that mPFC neuronal activity encoded relative positions between the start and goal. During awake immobility periods, the mPFC replayed temporally organized sequences of these generalized positions, resembling entire spatial trajectories. The occurrence of mPFC trajectory replay positively correlated with rule-switching performance. However, hippocampal and mPFC trajectory replay occurred independently, indicating different functions. These results demonstrate that the mPFC can replay ordered activity patterns representing generalized locations and suggest that mPFC replay might have a role in flexible behavior.

INTRODUCTION

Complex cognitive tasks usually incorporate multiple task stages and, as such, require the involvement of various brain areas. The brain areas involved in these tasks might coordinate their computations and work in concert or be required at different points in time and thus not show precise temporal coordination. Flexibly adapting the behavioral strategy following a change in a rule depends on the medial prefrontal cortex (mPFC). This dependency has been first described for humans and then also for rats, where mPFC lesions result in inflexibility and deficits in rule-switching performance (Floresco et al., 2008; Guise and Shapiro, 2017; Milner, 1963; Ragozzino et al., 1999). Spatial navigation and decision making, on the other hand, are functions largely attributed to the hippocampus (Morris et al., 1982). A task that combines the need for spatial navigation and the ability to flexibly shift strategies is the rule-switching task on the plus maze.

Spatially selective cells in the hippocampus, called place cells, are the substrate for the cognitive map required for spatial navigation (O'Keefe and Dostrovsky, 1971; O'Keefe and Nadel, 1978). Furthermore, the prominent role of the hippocampus in spatial functions has been partially attributed to replay, the reactivation of place cell activity patterns representing distant places or even recapitulating coherent trajectories (Foster and Wilson, 2006; Lee and Wilson, 2002; Nádasdy et al., 1999; Wilson and McNaughton, 1994). Such replay was first observed in sleep, where it is believed to play a role in memory consolidation (Wilson and McNaughton, 1994). Later, similar replay was seen during waking immobility periods, which, due to their temporal proximity to behavioral experience, suggested a function for replay in the solving of tasks. Indeed, awake replay in the hippocampus is thought to underlie the evaluation of past and future choices (Pfeiffer and Foster, 2013; Singer et al., 2013; Xu et al., 2019), and inhibiting hippocampal sharp wave ripples (SWRs) (periods of enhanced replay) in awake animals results in deficits in spatial task performance (Jadhav et al., 2012).

Similar to the hippocampus, replay in the mPFC was first observed in sleep, where task-induced sequential neural activity patterns become reactivated in subsequent rest periods (Euston et al., 2007). Furthermore, during sleep after a rule switch, there is more mPFC replay of waking neuronal activity patterns (Peyrache et al., 2009), raising the possibility that mPFC replay may strengthen knowledge for future rule switching. However, it is unclear so far whether the mPFC can replay sequences in the awake state and what information the reactivated neuronal patterns code for in rule switching.

Although the rule-switching task on the plus maze engages both the mPFC and hippocampus, it is not immediately obvious how replay in these two areas will relate to each other. Replay in the mPFC and hippocampus might occur at similar times, since the synchronization of oscillatory patterns increases between these two areas during rule switching (Benchenane et al., 2010). Moreover, while animals performed spatial tasks, the reactivation of correlated firing patterns of mPFC and hippocampal cell pairs has been demonstrated (Jadhav et al., 2016; Tang et al., 2017; Yu et al., 2018). However, mPFC and hippocampal replay in the plus-maze task might also occur independently, since the hippocampus, unlike the mPFC, has not been shown

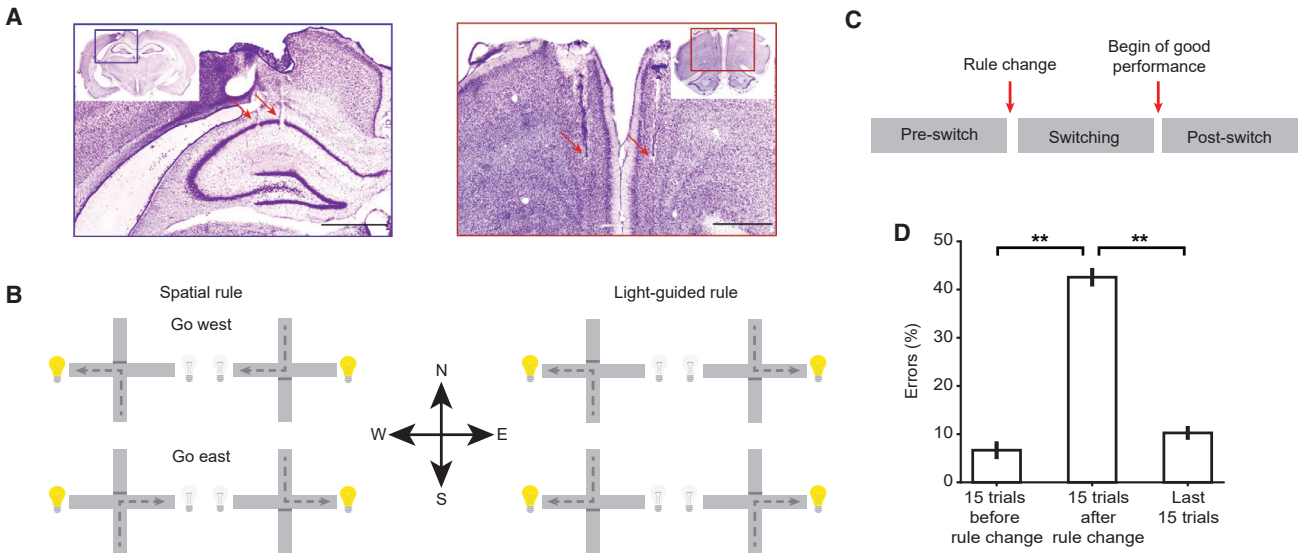

Figure 1. HPC and mPFC Neural Recordings during a Rule-Switching Task

(A) Tetrodes targeted the dorsal CA1 (unilaterally) and the prelimbic area of the mPFC (bilaterally). Inset shows an entire coronal brain slice with the box indicating the location of the magnified region. Red arrows indicate tetrode tips. Scale bar denotes 1 mm.

(B) Rats had to switch between a spatial- and light-guided rule on a plus maze. During the spatial rule, rats had to ignore the light cue, while in the light-guided task, rats had to follow it.

(C) Daily behavioral protocol comprising a pre-switch, switching, and post-switch block.

(D) Percentage of errors in 15 trials before rule change, 15 trials after rule change, and in the last 15 trials. A rule change lead to significantly more error trials (**$p < 0.01$, Wilcoxon signed-rank test). Error bars show standard error of the mean (SEM).

See also Figure S1.

to be required for rule switching. Therefore, replay in the mPFC during the rule-switching task might be involved in functions specific to that brain area.

Here, we examined the complementary roles of the hippocampus and mPFC in rule switching. The spatial coding of neuronal populations in these regions was compared. Moreover, we tested the role of mPFC replay in rule switching by investigating its content and relationship to behavior and by examining its coordination to hippocampal replay. In these analyses we specifically tested whether the mPFC can reactivate extended experiences, beyond cell pair measures, and with a behavioral relevance.

RESULTS

With 32-tetrode microdrives, we recorded from a total of 817 mPFC and 1,025 dorsal hippocampus CA1 (HPC) putative principal cells across 13 experimental sessions (averaging 63 mPFC and 78 HPC principal cells per session) (Figures 1A, S1A, and S1B). Four rats were trained on a rule-switching task on a plus maze where reward had to be collected following a spatial- or light-guided strategy (Figures 1B, S1C, and S1D). At the beginning of every trial, the animal was placed in one of the two start arms (north or south) to then approach the maze center and collect a reward in one of the two goal arms (east or west). On every recording day, the animal had to initially collect rewards based on the previous day's last rule (pre-switch) until the rule was changed unannounced. The animal had to then abandon the old rule, and through trial and error, switch to the new rule (switching) and successfully perform new rule trials (post-switch) (Figure 1C). Once the animal performed five consecutive correct trials after rule change, the post-switch block started from the third correct trial (i.e., beginning of good performance) and ended when a certain number of correct trials was reached. The unannounced rule change lead to animals making significantly more errors (15 trials before versus 15 trials after rule change: $p = 0.0014$, Wilcoxon signed-rank test; Figure 1D), but animals successfully updated to the new strategy in later trials (first 15 trials after rule change versus last 15 trials: $p = 0.0014$; Figure 1D).

Generalized Spatial Coding in the mPFC

We constructed spatial rate maps and, similar to previous studies, found that mPFC firing is related to locations in the arms of the plus maze (Fujisawa et al., 2008; Hok et al., 2005; Zielinski et al., 2019) (Figures 2A and S2A). However, it also became apparent that mPFC cells often had multiple firing fields and that these fields frequently occupied symmetrical locations on opposite arms of the maze, thereby generalizing between the two start or the two goal arms (Figure 2B). The symmetric coding properties of the mPFC were quantified with population vector correlations between the two start arms and between the two goal arms (Figure 2C). The average population vector correlation for the mPFC was 0.89 for the start arms and 0.92 for the goal arms, indicating that the mPFC population had similar spatial firing patterns in both start and goal arms.

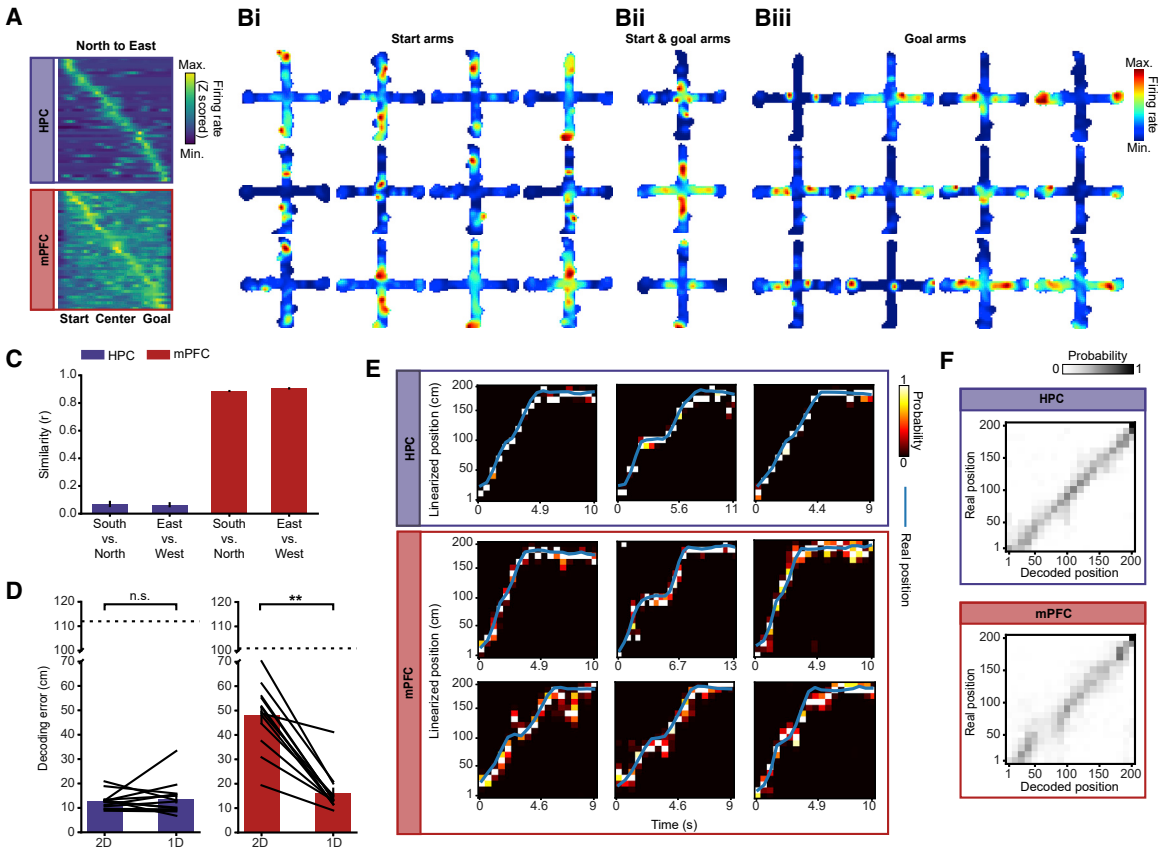

Figure 2. Coding of the Relative Spatial Position in the mPFC

(A) Linearized rate maps of one of the four trajectories (north to east) for HPC and mPFC cells that were simultaneously recorded during one session. Rate maps were normalized relative to the mean firing rate of the respective cell (*Z* scored) and sorted based on the position of a cell's peak firing rate.

(B) Rate maps from example mPFC cells illustrating their symmetric spatial firing and generalization between the two start and two goal arms. (i) Cells that fire in a similar position in both start arms. (ii) Cells that fire in a similar position in both start and both goal arms. (iii) Cells that fire in a similar position in both goal arms.

(C) Population vector correlations between the south and north and between the east and west arms quantifying their similarity. Error bars show SEM.

(D) Mean decoding error for Bayesian positional decoding performed with 2D or linearized 1D rate maps. For the mPFC, decoding performed with 1D rate maps led to a strong improvement in decoding precision. Every session is represented as a black line, and the horizontal dashed line indicates the chance level decoding error. Error bars show SEM (n.s., non significant, $^{**}p < 0.01$, Wilcoxon signed-rank test).

(E) 1D Bayesian decoding of HPC or mPFC spiking for example trials. For every time window of 0.512 ms, position probabilities are plotted with the animal's real position overlaid (blue line).

(F) Confusion matrix for the HPC and mPFC cell population, where the diagonal shows correspondence between the decoded and real positions.

See also Figure S2.

On the other hand, the two start and two goal arm firing patterns of the HPC ensemble were distinct (average population vector correlation, start arms: 0.07, goal arms: 0.06).

To further investigate whether the mPFC preferentially exhibits spatially discrete or rather generalized spatial firing, we applied the Bayesian method to decode the position of the animal during running epochs on the basis of mPFC spiking. We observed a strong improvement in decoding precision when the animal's relative position between start and goal was decoded (''1D decoding'') instead of requiring the decoding to distinguish between the different pairs of start and goal arms (''2D decoding''; mean decoding error: 2D = 46 ± 3 cm; 1D = 16 ± 3 cm; p = 0.0015, Wilcoxon signed-rank test; Figure 2D). 1D decoding was also performed based on HPC spiking; however, it yielded a decoding precision similar to that of 2D decoding (mean decoding error: 2D = 14 ± 1 cm; 1D = 14 ± 1 cm; p = 0.5, Wilcoxon signed-rank test; Figure 2D). Precision of 1D decoding performed on mPFC spiking was comparable to that performed on HPC spiking (p = 0.25, Wilcoxon signed-rank test), although the number of spikes required to decode the position with a mean decoding error of less than 30 cm was lower for the HPC (p = 0.0014, Wilcoxon signed-rank test; Figure S2B). 1D decoding of mPFC and HPC spiking corresponded with the real linearized position of the animal (Figures 2E and 2F).

The low 2D decoding performance for the mPFC did not result from the possibility that the mPFC differentiates between the correct and error arm; the mPFC population vectors constructed from only correct or error trials were highly similar (Figure S2Ci), and decoding the 2D location of correct trials with rate maps constructed from only correct trials did not improve decoding

performance ($p = 0.92$, Wilcoxon signed-rank test; Figure S2Cii). Furthermore, only for 59% of all trials could the trial outcome (i.e., correct or error) be decoded from mPFC spiking in the goal arm (Figure S2Ciii).

Similar to the HPC, mPFC 1D decoding error decreased with windows of increasing number of spikes, windows with increasing number of firing cells, and windows of increasing time lengths (Figure S2D). For the HPC and mPFC, the decoding errors in the pre-switch, switching, and post-switch blocks did not differ (HPC: $p = 0.16$, mPFC: $p = 0.63$, Kruskal-Wallis test; Figure S2E). It is important to note that although the mPFC predominantly coded for generalized positions, it did also show some allocentric spatial selectivity; 2D decoding with the mPFC population was nevertheless better than chance (Figure 2D, horizontal dashed line), and when using a Bayesian method to differentiate between the two goal arms, the goal arm identity could be distinguished with mPFC population firing at the goal arm in 76% of all trials (Figure S2F).

In summary, the improved precision with 1D decoding indicates that the mPFC, unlike the HPC, holds less information about discrete places. The highly similar firing patterns of the mPFC in the two start and the two goal arms suggests that the mPFC generalizes between places with similar task context. Due to the low mPFC 2D decoding performance and symmetric firing of mPFC cells, subsequent analyses were performed with 1D linearized rate maps where the two start and two goal arms are collapsed onto one linear map. For consistency, analyses on HPC spiking were also performed with linearized rate maps.

Non-local Positional Encoding in the mPFC

We next investigated whether the mPFC can activate non-local spatial locations (i.e., spatial locations not bound to the current position of the animal). Based on mPFC population firing, we decoded the position of the animal not only during running but also immobility periods (Figures S3D and S3E). Rate maps were computed using all trials except for the one on which positional decoding was performed and without speed filtering. This procedure ensured that any non-local position decoded from mPFC spiking recorded during immobility did not result from a lack of encoding information during immobility. We found that the encoded position can be highly non-local (Mashhoori et al., 2018), and that the majority of non-local events in the mPFC occurred during periods of low speed (Figure S3F). We observed two peaks in the distribution of decoding errors during immobility, with the highest peak at zero and a second, smaller one at ~150 cm away from the real location of the animal while immobile (Figure S3E). Also for the HPC, the majority of non-local events occurred during immobility (Figures S3A–S3C).

Replay of Trajectories in the mPFC

Hippocampal replay can recapitulate entire past or future trajectories that serves cognitive functions, such as memory consolidation and future planning (Csicsvari et al., 2007; Davidson et al., 2009; Diba and Buzsáki, 2007; Foster and Wilson, 2006). Experimental evidence demonstrated that extended replay is not limited to the hippocampus and can also be observed in other cortical areas (Ji and Wilson, 2007; O'Neill et al., 2017). Since the mPFC represented locations away from the animal's current position during immobility, we tested whether the mPFC links individual non-local positions to an ordered sequence of locations, thereby recapitulating an extended behavioral experience. For each session and separately for HPC and mPFC, we selected the number of spikes that allowed us to decode the position of the animal during running with an average error lower than 30 cm. Each immobility period was then divided into these spiking-windows that contained the same number of spikes and did not overlap (Stella et al., 2019). The distribution of spiking-window lengths was especially broad for the mPFC, indicating that compared to the HPC, mPFC population activity during replay is less synchronized and exhibits a weaker transient increase in activity (Figure S4A). Using the Bayesian position decoding algorithm and rate maps computed from spiking activity during running periods, a normalized likelihood function on the binned linearized maze positions was computed for each spiking-window. For the HPC and mPFC, the decoding confidence of a spiking-window (i.e., maximum likelihood divided by the sum of all likelihoods) only minimally, albeit significantly, decreased with spiking-window size, showing that even positions decoded from larger spiking-windows tend to point to similarly defined spatial locations (HPC: $p = 0.0016$, mPFC: $p = 0.0009$, Kruskal-Wallis test; Figure S4B). Then, for every four consecutive spiking-windows, we computed a trajectory score, which was based on the decoding probability between decoded positions (decoding probability score; ranging between 0 [least optimal] and 1 [most optimal]) and the smoothness of the trajectory (smoothness score; ranging between 0 [most optimal] and 1 [least optimal]). Both scores were tested against chance using a place field rotation shuffling procedure. We observed that the decoding probability score distributions were significantly different, with scores closer to zero for shuffled events (HPC and mPFC: $p < 0.00001$, two-sample Kolmogorov-Smirnov [KS] test; Figure 3A). The difference remained significant when taking nonoverlapping four-spiking-window events (HPC and mPFC: $p < 0.00001$, two-sample KS test; Figure S4C). The smoothness score distributions were also significantly different and were higher for the shuffled events (HPC and mPFC: $p < 0.00001$, two-sample KS test; Figure 3B). Also for the smoothness score, the difference remained significant when taking nonoverlapping four-spiking-window events (HPC and mPFC: $p < 0.00001$, two-sample KS test; Figure S4D). We unified the two measures by subtracting the smoothness score from the decoding probability score, yielding the trajectory score, a real number that ranges between -1 (least optimal) and 1 (most optimal). To identify significant trajectory events that were then used for subsequent analyses, we selected all the four-spiking-window events that had a trajectory score above the 95th percentile of their own shuffled distribution. Since replayed trajectories can be longer than four spiking-windows, we incrementally extended significant four-window events. As a result, 58% and 60% of the considered five-window events in the HPC and mPFC, respectively, passed their own shuffling (HPC and mPFC: $p < 0.00001$, binomial test). The same procedure was then performed with the extended event.

As a control, to verify the validity of longer than four-window events, we computed the scores for five- and six-spiking-window events, with overlapping windows as described above for

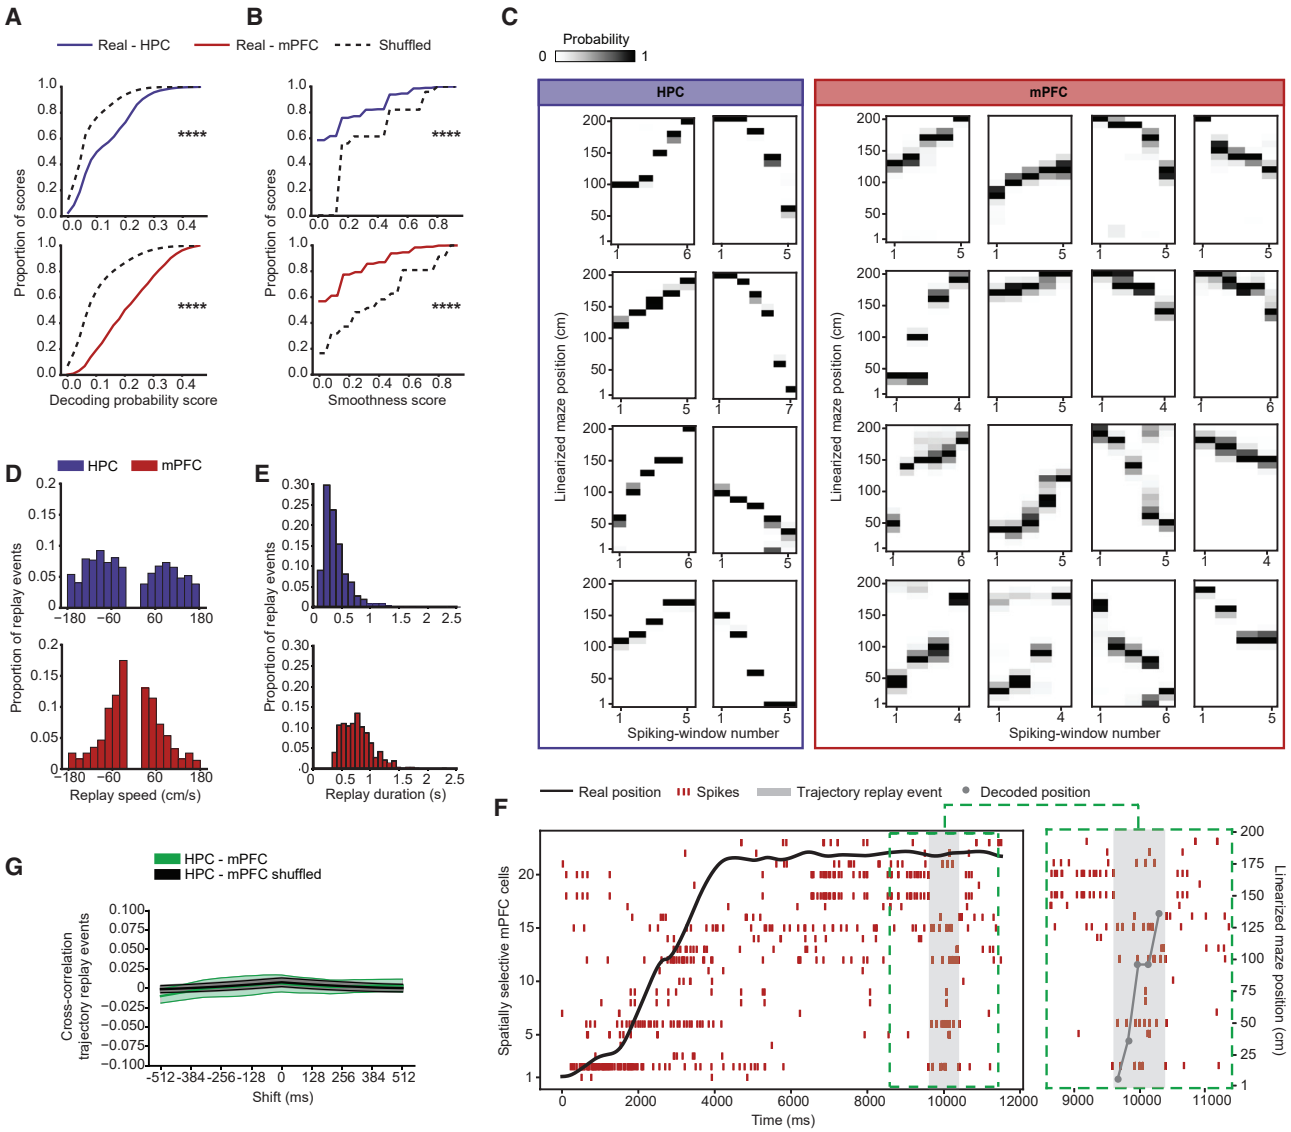

Figure 3. Replay of Trajectories in the mPFC during Immobility

Trajectory replay analysis was performed on windows with a fixed number of spikes (i.e., spiking-window), and the quality of linear replay was assessed with the decoding probability and smoothness scores.

(A) The cumulative distribution of decoding probability scores computed from real and place field-rotated shuffled data. For the HPC and mPFC, the decoding probability scores of the real data were significantly higher than of the shuffled data (****$p < 0.00001$, two-sample KS test).

(B) The cumulative distribution of smoothness scores computed from real and place field-rotated shuffled data. For the HPC and mPFC, the smoothness scores of the real data were significantly lower than of the shuffled data (****$p < 0.00001$, two-sample KS test).

(C) Examples of significant trajectory replay events in the HPC and mPFC for all four animals (rows). Decoding was performed on spiking-windows, and the position probabilities of a spiking-window were normalized by the maximum probability. Only significant trajectory events were used for subsequent analyses.

(D) Distribution of replaying speeds of trajectory replay events detected in the HPC and mPFC. One criterion for an event replaying a trajectory was a reactivation speed of above 20 cm/s. Positive and negative speeds indicate forward and reverse playing trajectories, respectively.

(E) Distribution of replay durations of trajectory replay events detected in the HPC and mPFC.

(F) Spiking of spatially selective mPFC cells (n = 23 cells, spatial information >0.1, sparsity <0.1) during an example trial. The black line denotes the real position of the animal and the gray shaded area an event where trajectory replay has been detected. The dashed box zooms into the trajectory replay event. Note the change in network activity (i.e., increase or decrease in firing rate of cells) during the trajectory replay event.

(G) Cross-correlation of trajectory replay events detected in the HPC and mPFC. The cross-correlation did not show any peaks and overlapped with that of the shuffled data, indicating that trajectory replay events in these two areas generally did not co-occur. Shaded areas show SEM.

See also Figures S3–S5.

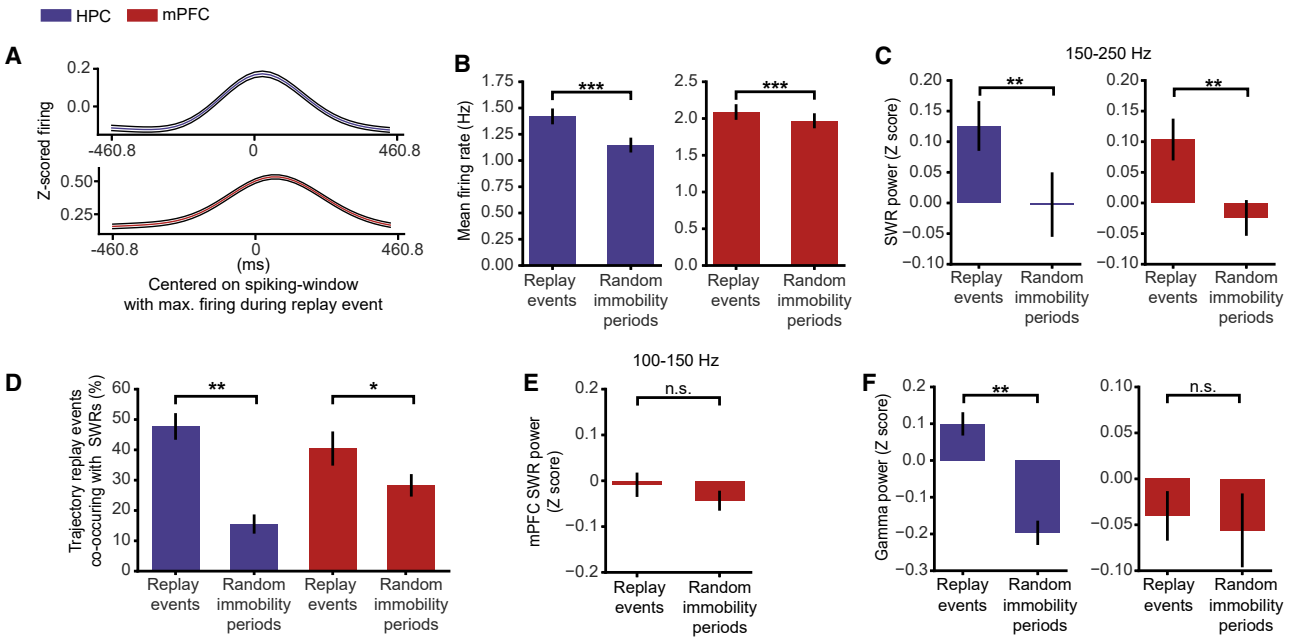

Figure 4. Network Activity during Trajectory Replay Events

(A) The population firing of HPC (top) and mPFC (bottom) cells during replay events in their respective area. The *Z*-scored population firing during the replay event was centered on the replay event spiking-window with the maximum firing rate. Black lines show SEM.

(B) The firing rate of HPC and mPFC cells increased during replay events in the respective area, relative to randomly selected immobility periods.

(C) The power in the hippocampal SWR band (150–250 Hz) was significantly higher during HPC and mPFC replay events compared to randomly selected immobility periods.

(D) The percentage of replay events that co-occurred with a hippocampal SWR event. SWRs were more often present during HPC and mPFC replay events than during randomly selected immobility periods.

(E) SWR power in the mPFC (100–150 Hz) was not different during mPFC replay events compared to randomly selected immobility events.

(F) Gamma (30–90 Hz) power during replay events. Hippocampal gamma power was increased during HPC replay events, while mPFC gamma power was not different during mPFC replay events. Error bars show SEM (n.s., non significant, *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$, Wilcoxon signed-rank test).

four spiking-windows, and observed that the distributions of scores were still significantly different from chance (decoding probability score: HPC and mPFC: all $p < 0.00001$; smoothness score: HPC and mPFC, all $p < 0.00001$; two-sample KS test; Figures S4E–S4H). Finally, when adding a fifth window to a significant four-window event, scores were also significantly different from those of the shuffled distribution in which only the fifth window was decoded using shuffled rate maps (decoding probability score: HPC and mPFC, $p < 0.00001$; smoothness score: HPC and mPFC: $p < 0.00001$, two-sample KS test; Figures S4I and S4J).

mPFC and HPC Trajectory Replay Events Occur Independently

As a result, 9% of the time spent immobile showed significant mPFC replay of trajectories. Selected mPFC events resembled trajectories typically decoded from the HPC, and both showed forward and reverse playing trajectories (Figures 3C and 3D). The distribution of trajectory speeds was wider in the HPC than mPFC (95th confidence intervals [cm/s]: HPC, [−491, 452]; mPFC, [−186, 191]; Figure 3D) and the duration of trajectory replay events in mPFC was longer than of those in the HPC (median: mPFC, 0.74 s; HPC, 0.33 s; $p < 0.00001$, Mann-Whitney *U* test; Figure 3E). Trajectory replay speed in both areas was associated with increased population firing rate in the same area (HPC and mPFC: $p < 0.0001$, one-sample t test; Figure S5A). During mPFC trajectory replay events, mPFC cells start or stop firing, similar to the sudden change in network activity typically seen for HPC replay events (Figure 3F). Finally, cross-correlating the occurrence of significant trajectory replay events in the HPC with those in the mPFC showed that trajectory replay in these two areas was independent of each other (Figure 3G). On average, only 5% of all trajectory events in the mPFC occurred simultaneously with trajectory events in the HPC.

Robustness of mPFC Trajectory Replay Events

In addition to the decoding probability and smoothness scores, we used further measures to detect significant trajectory replay events (Davidson et al., 2009; Gupta et al., 2010). Significant trajectory events that were detected with the methods described above were taken, and, with the approaches from the Davidson et al. and Gupta et al. studies, scores were computed over these events. For both methods, replay events in the HPC and mPFC had larger scores than the shuffled data (all $p < 0.00001$, two-sample KS test; Figures S5B and S5C). This demonstrates that awake trajectory replay in the mPFC can even be detected with other replay quality measures. Furthermore, decoding positions with rate maps of smaller bin sizes (i.e., 5 cm and

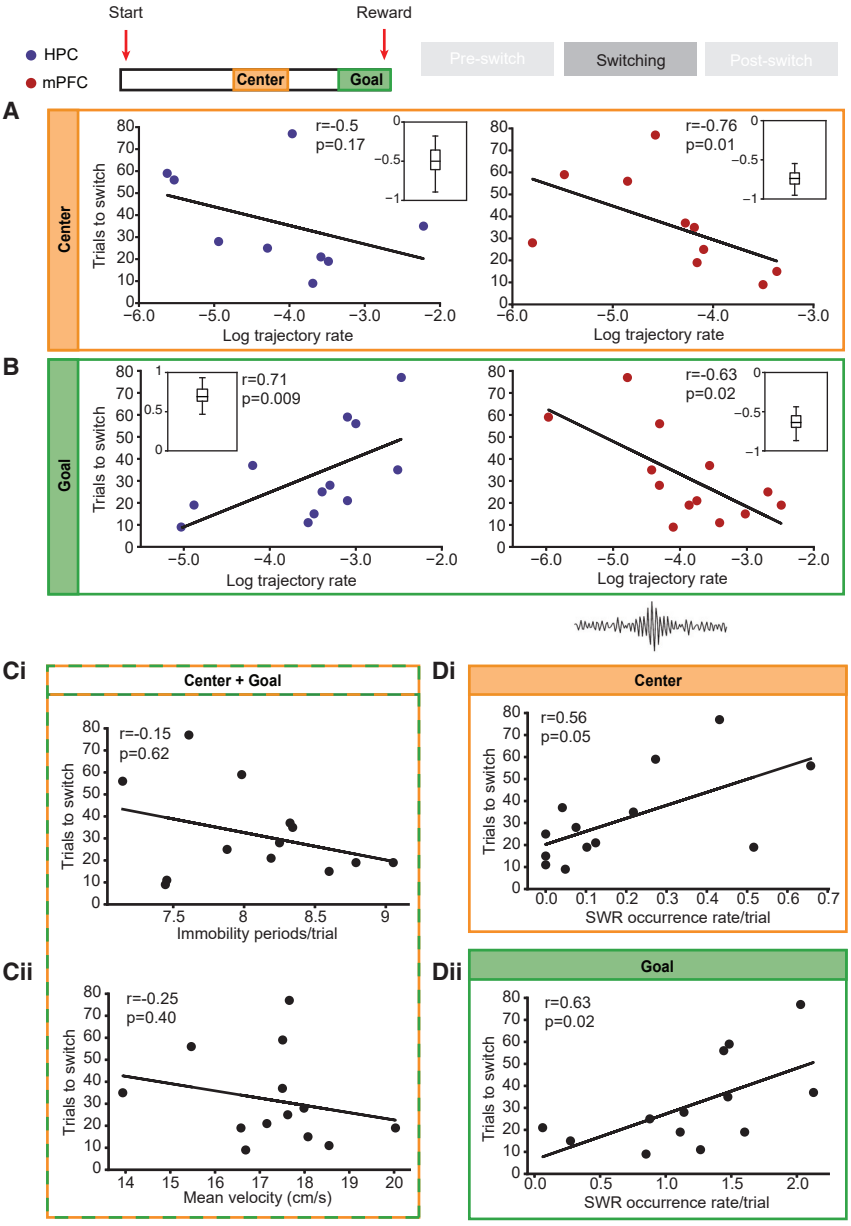

Figure 5. mPFC and HPC Trajectory Replay Oppositely Correlate with Rule-Switching Performance

(A) For every session, a mean trajectory replay rate was calculated over switching block trials. mPFC trajectory replay rate at the maze center negatively correlated with number of trials required to switch to the new rule and hence positively correlated with rule-switching performance. Each dot denotes one session and sessions with no trajectory replay for a condition were excluded. Inset shows mean r value obtained from 100 correlations of bootstrapped data.

(B) At the goal, the trajectory replay rate in the HPC positively correlated with number of trials required to switch to the new rule and therefore negatively correlated with rule-switching performance. Conversely, also at the goal, the mPFC trajectory replay rate negatively correlated with number of trials required to switch.

(C) Correlating number of trials to switch with other aspects of behavior. (i) The number of immobility periods per trial was not correlated with the number of trials required to switch. (ii) The mean running velocity in a trial was not correlated with the number of trials required to switch.

(D) SWR events were detected during immobility periods and the SWR number in a trial normalized by the immobility during that particular trial. The mean SWR event rate per trial was computed. (i) The SWR event rate at the center showed a trend toward significantly correlating with number of trials required to switch ($p = 0.05$). (ii) The SWR event rate at the goal positively correlated with the number of trials required to switch. (Spearman's rank-order correlation was used for all correlations here.)

See also Figure S6.

2.5 cm) still resulted in events that had higher trajectory scores than shuffled (all $p < 0.00001$, two-sample KS test; Figures S5D and S5E). Finally, since previous replay studies used time windows of a fixed length instead of the variable-length spiking-window used here, we computed the trajectory scores with windows of fixed time. Using fixed time windows resulted in trajectory scores that were still significantly different from scores of the shuffled distribution (all $p < 0.00001$, two-sample KS test; Figure S5F).

We then constructed separate rate maps from spiking during running and during immobility periods with trajectory replay and without. Correlating these rate maps for every mPFC cell showed that the running and immobility with trajectory replay maps were significantly less similar than the running and immobility without replay maps ($p < 0.00001$, two-sample KS test). Furthermore, the firing rates of mPFC cells during running were positively correlated with the firing during immobility with trajectory replay ($r = 0.90$, $p < 0.0001$, Spearman rank-order correlation). We therefore also observed mPFC replay on a single-cell level, since during replay, the same cells were active as during running periods (Pavlides and Winson, 1989) and because the firing location of cells could decouple from the current location of the animal.

Trajectory Replays in the mPFC Co-occur with HPC SWRs

We investigated whether replay events were associated with population-level responses. The population firing rate of mPFC neurons peaked during mPFC replay events (Figure 4A), similar to HPC population responses during HPC replay (Figure 4A). Then, for every replay event, a random immobility period of the same length was selected and the average firing rate computed for each cell during replay and random immobility periods. For both the HPC and mPFC, the firing rate of cells increased during replay

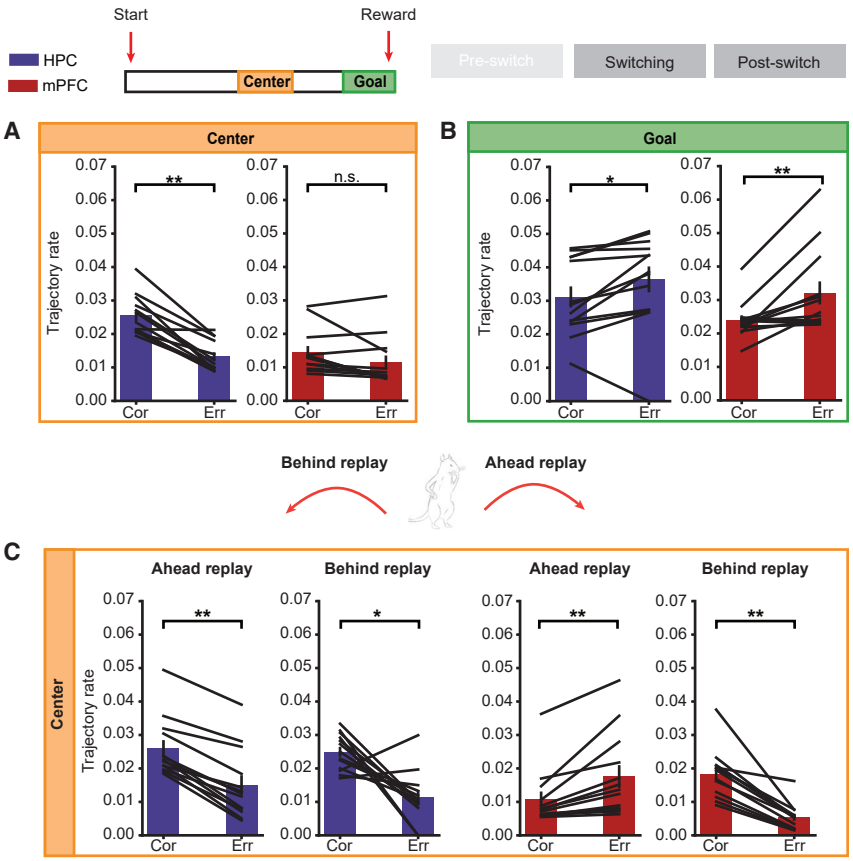

Figure 6. A Role for mPFC Trajectory Replay during Error Trials in Switching Performance

A mean trajectory replay rate was calculated over switching and post-switch block trials.

(A) HPC, but not mPFC, trajectory replay rate in the center was higher during correct trials than error trials.

(B) Trajectory replay rate at the goal was higher during error trials than correct trials for both the HPC and mPFC.

(C) Separating trajectory replay in the center into ending ahead (thus pointing toward the goal) and ending behind (pointing toward the start) the animal, showed that for the mPFC, ahead replay was increased during error trials, while behind replay was increased during correct trials. Ahead and behind trajectory replay in the HPC did not seem to support different functions, and both were increased during correct trials. Error bars show SEM (n.s., non significant, *$p < 0.05$, **$p < 0.01$, Wilcoxon signed-rank test, Bonferroni-Holm correction).

events compared to the randomly selected baseline firing rate (HPC and mPFC: $p < 0.001$, Wilcoxon signed-rank test; Figure 4B).

We then examined whether our detected trajectory replay events co-occurred with awake hippocampal SWRs. Hippocampal SWR power was increased during HPC replay events ($p = 0.005$, Wilcoxon signed-rank test; Figure 4C), and HPC replay events co-occurred more often with hippocampal SWRs than randomly selected immobility periods ($p = 0.001$, Wilcoxon signed-rank test; Figure 4D). It is possible that low-power SWRs were not detected, which could explain why not all HPC replay events coincided with SWRs. We performed the same analysis for mPFC replay events and also found an increase in hippocampal SWR power ($p = 0.001$, Wilcoxon signed-rank test; Figure 4C). Furthermore, 40% of mPFC replay events co-occurred with hippocampal SWR events, which was significantly higher than the co-occurrence of randomly selected immobility events and SWRs ($p = 0.02$, Wilcoxon signed-rank test; Figure 4D). Therefore, although mPFC replay events were temporarily uncorrelated with our detected replay events in the HPC, they nonetheless could co-occur with hippocampal SWR activity. Considering that half of our HPC replay events did not co-occur with SWRs and mPFC replay was generally associated with SWRs, we cross-correlated the occurrence of mPFC trajectory replay events with those HPC events that co-occurred with SWRs. However, also here, the cross-correlation did not indicate that trajectory events in the HPC and mPFC are dependent ($p = 1$, Kruskal-Wallis test).

It has been reported that cortical areas can also exhibit ripple frequency oscillations (Khodagholy et al., 2017). However, we did not find differences in mPFC ripple power (100–150 Hz) during replay and random immobility periods ($p = 0.12$, Wilcoxon signed-rank test; Figure 4E). Finally, the power of hippocampal gamma frequency oscillations (30–90 Hz) was slightly increased during HPC replay events, while gamma power in the mPFC did not change during mPFC replay events (HPC: $p = 0.002$, mPFC: $p = 0.7$, Wilcoxon signed-rank test; Figure 4F).

Trajectory Replay in the mPFC Correlates with Performance

There is evidence that awake replay in the hippocampus supports spatial task performance (Jadhav et al., 2012; Singer et al., 2013). To investigate whether awake trajectory replay in the mPFC goes beyond a mere recapitulation of experience and possibly supports the mPFC's role in rule switching, we studied the relationship of awake trajectory replay with rule-switching performance. For every trial, the number of trajectory replay events was divided by the immobility time to obtain the rate of trajectory replay. Then, the mean trajectory rate in the switching block was plotted against the number of trials required to switch to the new rule. The trajectory replay rate in the maze center for the HPC did not correlate with performance ($r = -0.5$, $p = 0.17$, Spearman rank-order correlation; Figure 5A). However, the rate of mPFC trajectory replay events, at both the center and goal, negatively correlated with the number of trials required to switch to the new rule and hence correlated positively with performance (center: $r = -0.76$, $p = 0.01$, goal: $r = -0.63$, $p = 0.02$, Spearman rank-order correlation; Figures 5A and 5B). Conversely, HPC trajectory replay in the goal area positively correlated with trials to switch and hence showed a negative correlation with performance ($r = 0.71$, $p = 0.009$, Spearman rank-

order correlation; Figure 5B). In light of the correlation between mPFC replay and switching performance, we investigated whether the occurrence of a mPFC replay event at the goal resulted in a higher likelihood of the next trial to be correct. However, the probability of the next trial to be correct did not differ depending on whether there was mPFC replay at the goal or not ($p = 0.92$, Wilcoxon signed-rank test; Figure S6A).

When significant trajectory events had to pass an additional cell identity shuffling procedure, 85% and 87% of HPC and mPFC events, respectively, were still significant (Figure S6B). Furthermore, all the correlations between trajectory replay rate and switching performance were maintained (Figures S6C and S6D). We tested whether other behavioral aspects could explain the correlation results. However, the number of trials required to switch was not correlated with the number of immobility periods per trial or the mean movement speed of the animal (immobility periods: $r = -0.15$, $p = 0.62$, mean velocity: $r = -0.25$, $p = 0.40$, Spearman rank-order correlation; Figure 5C). Interestingly, there was a positive correlation between the number of SWR events per trial at the goal and trials required to switch ($r = 0.63$, $p = 0.02$, Spearman rank-order correlation; Figure 5D). This result is in agreement with the positive correlation between HPC trajectory replay rate at the goal and the number of trials required to switch (Figure 5B).

Trajectory Replay and Trial Outcome

We then computed the mean trajectory rate separately for correct and error trials over all trials after the rule switch (i.e., trials from the switching and post-switch block). The mPFC trajectory replay event rate at the goal was higher in error than correct trials ($p = 0.009$, Wilcoxon signed-rank test; Figure 6B). This result and the positive correlation with performance provide support that mPFC trajectory replay at the goal area during error trials has a facilitating role in switching to the new rule. For the HPC, goal trajectory replay was slightly increased in error trials ($p = 0.016$, Wilcoxon signed-rank test; Figure 6B). HPC trajectory replay in the center was higher during correct than in error trials, also when separating for trajectory replay ending ahead and behind the animal (total: $p = 0.0015$, ahead: $p = 0.006$, behind: $p = 0.02$, Wilcoxon signed-rank test, Bonferroni-Holm correction; Figures 6A and 6C). This is in agreement with previous reports that the hippocampus shows more coordinated replay before correct trials than before error trials (Singer et al., 2013). Separating ahead and behind trajectory replays in the center for the mPFC showed that while there were more ahead trajectories during error trials, there were more behind trajectories during correct trials (total: $p = 0.05$, ahead: $p = 0.006$, behind: $p = 0.003$, Wilcoxon signed-rank test, Bonferroni-Holm correction; Figures 6A and 6C), suggesting that ahead and behind trajectories in the mPFC might serve different functions.

Although we detected replay events with linearized rate maps that do not differentiate between the two start or the two goal arms, we nonetheless attempted to determine whether trajectory replay events at the center and goal preferentially replayed the current or previous goal arm. After all, with our mPFC data we could differentiate between the two goal arms with 76% accuracy (Figure S2F). At the center, replay events ending ahead of the animal of neither brain area were biased toward the past or upcoming goal arm (HPC, previous: $p = 0.99$, current: $p = 0.36$; mPFC, previous: $p = 0.98$, current: $p = 0.5$, binomial test; Figure S6E). At the goal, replay events in the HPC and mPFC showed a strong preference for the current goal arm (HPC, previous: $p < 0.0001$, current: $p < 0.00001$; mPFC, previous: $p = 0.004$, current: $p < 0.00001$, binomial test; Figure S6F).

DISCUSSION

A number of past studies found that mPFC neurons display place-related firing (Fujisawa et al., 2008; Hok et al., 2005; Zielinski et al., 2019). We too found that mPFC neurons can hold spatial information. However, in our case, mPFC cells primarily represented a generalized form of space; instead of coding for specific locations, mPFC neurons encoded the relative position between the start and the goal. The mPFC might therefore have a role in the generalization of places with similar meaning (Yu et al., 2018). The symmetric shape of the maze used in this study and the similar task-related meaning of the two start and two goal arms might have made the symmetric firing properties of mPFC cells more evident than other tasks or mazes would have. At the population level, the mPFC could also encode goal arm identity. However, goal arm decoding accuracy was only 76%, and population vector analysis indicated that mPFC assembly patterns were very similar in the two goal arms. Therefore, either a small, independent population of mPFC cells held information about the different goals or cells signaled goal identity with small firing rate deviations. Nevertheless, two independent coding mechanisms were present in the mPFC encoding relative position and goal identity.

We found that during awake immobility, the mPFC can reactivate places distant from the actual position of the animal. This is in line with a previous study reporting that the mPFC can recall reward locations that are remote from the animal, resembling replay in the hippocampus (Mashhoori et al., 2018). However, we demonstrated that awake mPFC replay can go beyond reactivation of single non-local positions. The mPFC replayed temporally organized sequences of positions recapitulating generalized behavioral trajectories. Such replay occurred not only during rule switching but also during stable performance of the spatial or cue-guided rule on the plus maze. The mPFC trajectory replay that we detected progressed not only in the forward order but also in the reverse order, indicating that it is not a mere recapitulation but an abstraction of the original experience. Indeed, our findings on the relationship of mPFC replay with rule-switching performance and its increase during error trials suggest that temporally organized replay in the mPFC may be actively involved in rule switching.

What might be the role of mPFC trajectory replay during awake immobility in the rule-switching task? To address this question, one has to keep in mind that, according to our data, the mPFC primarily retains only part of the spatial information found in hippocampal representations; it produces a trajectory-independent representation by generalizing between the two start and two goal arms. Therefore, since the fundamental elements comprising the task entail running down a start arm and, at the maze center, choosing one of the goal arms, the mPFC might

tease out the elements common to both rules. The mPFC holds the neurobiological basis of schemas that are higher-level frameworks of knowledge onto which new, related information can be rapidly assimilated (Tse et al., 2007). In the present task, the general task structure may be represented in the mPFC as a schema onto which the new rule can be mapped. We found that the occurrence of mPFC trajectory replay at the goal and those trajectory replays at the center pointing toward the goal were increased on error trials. Moreover, before the animal switched to the new rule, the overall occurrence rate of mPFC replay positively correlated with how fast the animal switched. Therefore, during periods before the animal adapts to the new rule, replay of the general task structure might aid the determination of the specific rule currently in place. Alternatively, since positions of the current goal arm were reactivated in a significant fraction of mPFC replay events at the goal, mPFC trajectory replay may constitute an evaluation of the choice and trial outcome (Narayanan and Laubach, 2008; Passecker et al., 2019; Sul et al., 2010). A single mPFC replay event at the goal, however, did not immediately result in better performance on the next trial, indicating that mPFC replay may rather be part of an evidence accumulating process spanning several trials and not reflective of a moment of sudden insight (Durstewitz et al., 2010). That goal replay in the HPC negatively correlated with switching performance and was increased during error trials suggests that HPC goal replay may only signal errors and does not have a direct role in facilitating rule switching.

In the present study, trajectory replay rate of the HPC and mPFC at the goal were oppositely correlated with rule-switching performance, and HPC and mPFC goal-directed (i.e., ending ahead of the animal) trajectory replays at the center oppositely correlated with error trials. Moreover, we did not find a temporal correlation between trajectory replay events in the HPC and mPFC. This was unexpected, since mPFC replay so far was mostly studied in temporal alignment with hippocampal SWRs, periods of enhanced hippocampal replay (Jadhav et al., 2016; Peyrache et al., 2009; Shin et al., 2019). However, replay in areas outside the hippocampus does not necessarily have to be linked with replay in the hippocampus. For example, independent replay has also been reported for the entorhinal cortex, a structure that is neighboring the hippocampus (O'Neill et al., 2017). It is therefore possible that mPFC replay serves different functions depending on whether it occurs independently or together with hippocampal replay. Synchrony between hippocampal and cortical replay may thereby have a consolidative role (Diekelmann and Born, 2010; Ji and Wilson, 2007), whereas independent replay may serve computational processes within the mPFC. Our findings therefore point to a complementary role of trajectory replay occurring in the mPFC and hippocampus during rule switching and in the subsequent maintenance of updated rules. Nonetheless, since mPFC trajectory replay significantly co-occurred with hippocampal SWRs, it may still be influenced by hippocampal activity. Furthermore, there is also the possibility that structures upstream of the mPFC, such as the medial entorhinal cortex, which can exhibit replay that is independent from that of the hippocampus (O'Neill et al., 2017), have an influence on mPFC replay.

Altogether, the presence of sequential reactivation of locations in the mPFC during the awake state indicates that temporally organized replay is a neural computation mechanism common to different brain areas (Ji and Wilson, 2007; Ólafsdóttir et al., 2016; O'Neill et al., 2017). When a higher-order organization was found for hippocampal replay, it opened new possibilities in which hippocampal mnemonic computations could be analyzed. Similarly, the discovery of awake trajectory replay in the mPFC and its relevance to the performance of rule switching might offer new ways in which mPFC function can be investigated.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- LEAD CONTACT AND MATERIALS AVAILABILITY
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Surgery
  - Plus maze Apparatus and Task
  - Histology and Reconstruction of Recording Positions
  - Data Acquisition
  - Spike Sorting and Unit Classification
- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Behavior
  - Linearized Position
  - Linearized Firing Rate Maps
  - Population Vector Similarity
  - Bayesian Decoding
  - Confusion Matrix
  - Cross-validation
  - Decoding Windows
  - Decoding of Trial Outcome and Goal Arm Identity
  - Trajectory Score
  - Shuffling Procedure
  - Trajectory Selection and Concatenation
  - Cross-correlation
  - Additional Quantification of Linear Replay Quality
  - Decoding of Goal Arms within Replay Events
  - Population Firing Rate during Replay
  - Oscillation Analyses
  - Behavioral Correlates
  - Statistical Analysis
- DATA AND CODE AVAILABILITY

SUPPLEMENTAL INFORMATION

Supplemental Information can be found online at https://doi.org/10.1016/j.neuron.2020.01.015.

A video abstract is available at https://doi.org/10.1016/j.neuron.2020.01.015#mmc3.

ACKNOWLEDGMENTS

We thank Todor Asenov and Thomas Menner from the Machine Shop for the drive design and production, Hugo Malagon-Vina for assistance in maze

automatization, Jago Wallenschus for taking the images of the histology, and Federico Stella and Juan Felipe Ramirez-Villegas for comments on an earlier version of the manuscript. This work was supported by the EU-FP7 MC-ITN IN-SENS (grant 607616).

AUTHOR CONTRIBUTIONS

K.K. and J.C. designed the study. K.K. and K.B. performed surgeries. K.K. executed the experiments. K.K. and M.N. analyzed the data. K.K., M.N., and J.C. wrote the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: July 30, 2019
Revised: December 19, 2019
Accepted: January 13, 2020
Published: February 6, 2020

REFERENCES

Benchenane, K., Peyrache, A., Khamassi, M., Tierney, P.L., Gioanni, Y., Battaglia, F.P., and Wiener, S.I. (2010). Coherent theta oscillations and reorganization of spike timing in the hippocampal-prefrontal network upon learning. Neuron *66*, 921–936.

Csicsvari, J., Hirase, H., Czurko, A., and Buzsáki, G. (1998). Reliability and state dependence of pyramidal cell-interneuron synapses in the hippocampus: an ensemble approach in the behaving rat. Neuron *21*, 179–189.

Csicsvari, J., O'Neill, J., Allen, K., and Senior, T. (2007). Place-selective firing contributes to the reverse-order reactivation of CA1 pyramidal cells during sharp waves in open-field exploration. Eur. J. Neurosci. *26*, 704–716.

Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). Hippocampal replay of extended experience. Neuron *63*, 497–507.

Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. Nat. Neurosci. *10*, 1241–1242.

Diekelmann, S., and Born, J. (2010). The memory function of sleep. Nat. Rev. Neurosci. *11*, 114–126.

Durstewitz, D., Vittoz, N.M., Floresco, S.B., and Seamans, J.K. (2010). Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. Neuron *66*, 438–448.

Euston, D.R., Tatsuno, M., and McNaughton, B.L. (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. Science *318*, 1147–1150.

Floresco, S.B., Block, A.E., and Tse, M.T.L. (2008). Inactivation of the medial prefrontal cortex of the rat impairs strategy set-shifting, but not reversal learning, using a novel, automated procedure. Behav. Brain Res. *190*, 85–96.

Foster, D.J., and Wilson, M.A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. Nature *440*, 680–683.

Fujisawa, S., Amarasingham, A., Harrison, M.T., and Buzsáki, G. (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. Nat. Neurosci. *11*, 823–833.

Guise, K.G., and Shapiro, M.L. (2017). Medial prefrontal cortex reduces memory interference by modifying hippocampal encoding. Neuron *94*, 183–192.e8.

Gupta, A.S., van der Meer, M.A.A., Touretzky, D.S., and Redish, A.D. (2010). Hippocampal replay is not a simple function of experience. Neuron *65*, 695–705.

Harris, K.D., Henze, D.A., Csicsvari, J., Hirase, H., and Buzsáki, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. J. Neurophysiol. *84*, 401–414.

Hok, V., Save, E., Lenck-Santini, P.P., and Poucet, B. (2005). Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. Proc. Natl. Acad. Sci. USA *102*, 4602–4607.

Jadhav, S.P., Kemere, C., German, P.W., and Frank, L.M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. Science *336*, 1454–1458.

Jadhav, S.P., Rothschild, G., Roumis, D.K., and Frank, L.M. (2016). Coordinated excitation and inhibition of prefrontal ensembles during awake hippocampal sharp-wave ripple events. Neuron *90*, 113–127.

Ji, D., and Wilson, M.A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. Nat. Neurosci. *10*, 100–107.

Khodagholy, D., Gelinas, J.N., and Buzsáki, G. (2017). Learning-enhanced coupling between ripple oscillations in association cortices and hippocampus. Science *358*, 369–372.

Lee, A.K., and Wilson, M.A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. Neuron *36*, 1183–1194.

Mashhoori, A., Hashemnia, S., McNaughton, B.L., Euston, D.R., and Gruber, A.J. (2018). Rat anterior cingulate cortex recalls features of remote reward locations after disfavoured reinforcements. eLife *7*, e29793.

Milner, B. (1963). Effects of different brain lesions on card sorting: the role of the frontal lobes. Arch. Neurol. *9*, 90–100.

Morris, R.G., Garrud, P., Rawlins, J.N., and O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. Nature *297*, 681–683.

Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., and Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. J. Neurosci. *19*, 9497–9507.

Narayanan, N.S., and Laubach, M. (2008). Neuronal correlates of post-error slowing in the rat dorsomedial prefrontal cortex. J. Neurophysiol. *100*, 520–525.

O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res. *34*, 171–175.

O'Keefe, J., and Nadel, L. (1978). The Hippocampus as a Cognitive Map (Oxford University Press).

O'Neill, J., Boccara, C.N., Stella, F., Schoenenberger, P., and Csicsvari, J. (2017). Superficial layers of the medial entorhinal cortex replay independently of the hippocampus. Science *355*, 184–188.

Ólafsdóttir, H.F., Carpenter, F., and Barry, C. (2016). Coordinated grid and place cell replay during rest. Nat. Neurosci. *19*, 792–794.

Passecker, J., Mikus, N., Malagon-Vina, H., Anner, P., Dimidschstein, J., Fishell, G., Dorffner, G., and Klausberger, T. (2019). Activity of prefrontal neurons predict future choices during gambling. Neuron *101*, 152–164.e7.

Pavlides, C., and Winson, J. (1989). Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. J. Neurosci. *9*, 2907–2918.

Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S.I., and Battaglia, F.P. (2009). Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. Nat. Neurosci. *12*, 919–926.

Pfeiffer, B.E., and Foster, D.J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. Nature *497*, 74–79.

Ragozzino, M.E., Detrick, S., and Kesner, R.P. (1999). Involvement of the prelimbic-infralimbic areas of the rodent prefrontal cortex in behavioral flexibility for place and response learning. J. Neurosci. *19*, 4585–4594.

Shin, J.D., Tang, W., and Jadhav, S.P. (2019). Dynamics of awake hippocampal-prefrontal replay for spatial learning and memory-guided decision making. Neuron *104*, 1110–1125.e7.

Singer, A.C., Carr, M.F., Karlsson, M.P., and Frank, L.M. (2013). Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. Neuron *77*, 1163–1173.

Stella, F., Baracskay, P., O'Neill, J., and Csicsvari, J. (2019). Hippocampal reactivation of random trajectories resembling Brownian diffusion. Neuron *102*, 450–461.e7.

Sul, J.H., Kim, H., Huh, N., Lee, D., and Jung, M.W. (2010). Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. Neuron *66*, 449–460.

Tang, W., Shin, J.D., Frank, L.M., and Jadhav, S.P. (2017). Hippocampal-prefrontal reactivation during learning is stronger in awake compared with sleep states. J. Neurosci. *37*, 11789–11805.

Tse, D., Langston, R.F., Kakeyama, M., Bethus, I., Spooner, P.A., Wood, E.R., Witter, M.P., and Morris, R.G.M. (2007). Schemas and memory consolidation. Science *316*, 76–82.

Wilson, M.A., and McNaughton, B.L. (1994). Reactivation of hippocampal ensemble memories during sleep. Science *265*, 676–679.

Xu, H., Baracskay, P., O'Neill, J., and Csicsvari, J. (2019). Assembly responses of hippocampal CA1 place cells predict learned behavior in goal-directed spatial tasks on the radial eight-arm maze. Neuron *101*, 119–132.e4.

Yu, J.Y., Liu, D.F., Loback, A., Grossrubatscher, I., and Frank, L.M. (2018). Specific hippocampal representations are linked to generalized cortical representations in memory. Nat. Commun. *9*, 2209.

Zhang, K., Ginzburg, I., McNaughton, B.L., and Sejnowski, T.J. (1998). Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. J. Neurophysiol. *79*, 1017–1044.

Zielinski, M.C., Shin, J.D., and Jadhav, S.P. (2019). Coherent coding of spatial position mediated by theta oscillations in the hippocampus and prefrontal cortex. J. Neurosci. *39*, 4550–4565.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental Models: Organisms/Strains		
Male Long-Evans rats	Janvier, France	RjOrl:LE
Software and Algorithms		
KlustaKwik	Harris et al., 2000	http://klustakwik.sourceforge.net
Python	Python	https://www.python.org
Other		
12 μm tungsten wires	California Fine Wire	Cat # CFW0010954
Headstage amplifier	Axona, St. Albans, UK	http://www.axona.com/

LEAD CONTACT AND MATERIALS AVAILABILITY

Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Jozsef Csicsvari (jozsef.csicsvari@ist.ac.at). This study did not generate new unique reagents.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Four male Long-Evans rats (300-350 g, 2-4 months of age; Janvier, France) were used in this study. The animals were housed in a separate room on a 12 hour light/dark cycle and were taken to the recording room each day prior to the experiments. Animals shared a cage with littermates before surgery. All procedures involving experimental animals were carried out in accordance with Austrian animal law (Austrian federal law for experiments with live animals) under a project license approved by the Austrian Federal Science Ministry (License number: BMWFW-66.018/0015-WF/V3b/2014).

METHOD DETAILS

Surgery

Rats were implanted with microdrives housing 32 individually-movable tetrodes, arranged into three bundles targeting the right dorsal hippocampus (specifically dorsal CA1, HPC) and left and right medial prefrontal cortex (specifically prelimbic area, mPFC). The HPC bundle consisted of 16 tetrodes and the two mPFC bundles of 8 tetrodes each. Tetrodes were fabricated out of four 12 μm tungsten wires (California Fine Wire Company, Grover Beach, CA) that were twisted and then heated to bind into a single bundle. Tetrode bundle lengths were cut so that the two mPFC bundles were 1-1.5 mm longer than the HPC bundle. The tips of the tetrodes were gold-plated to reduce the impedance to around 300 kΩ. Before surgery the animal was put under deep anesthesia using isoflurane (0.5%–3%), oxygen (1–2 L/min), and an initial injection of buprenorphine (0.1 mg/kg) and ketamine/xylazine (7:3 ketamine (10%) and xylazine (2%), 0.05ml/100 g). Craniotomies were drilled above the HPC (AP: −2.50 to −4.50, ML: −1.2 to −3.6) and above the mPFC across the sinus (AP: 4.60 to 2.50, ML: 0 to ± 0.8). Six anchoring screws were fixed onto the skull and two ground screws were positioned above the cerebellum. After dura removal the tetrode bundles we centered above their respective craniotomies and lowered into the brain at a depth of ∼2 mm for the mPFC and ∼1 mm for the HPC. The exact depth of mPFC tetrode implantation was noted to ensure later lowering into the target area. Tetrodes and craniotomies were coated in paraffin wax and the microdrive was anchored to the skull and screws with dental cement. The analgesic meloxicam (5 mg/kg) was given up to three days after surgery and the animal was allowed one week recovery. Thereafter, tetrodes were gradually moved in 50-200 μm steps into the HPC pyramidal cell layer and mPFC.

Plus maze Apparatus and Task

Following the recovery period, animals were food-restricted with *ad libitum* access to water and accustomed to the plus maze and rest box. The plus maze was elevated (80 cm) and consisted of four arms (85 cm long and 12 cm wide), referred to as north, east, south and west, and a connecting center. The animal was placed in one of the two start arms (north or south) and had to collect a food reward (MLab rodent tablet 20mg, TestDiet, Richmod, USA) in one of the two goal arms (east or west), depending on the rule employed. Access to the arm not chosen as the start was restricted, so that the maze became T-shaped. A small light at the end of one of the two goal arms was switched on. Which arm was chosen as the start and light-on arm was chosen pseudorandomly

for every trial, ensuring that an arm was not chosen more than three consecutive times. Once the animal reached a goal arm and ~5 s passed, the animal was manually picked up and placed in the rest box before commencing to the next trial after a delay of ~10 s. The animal had to retrieve the reward based on a spatial or response (light) rule. During the spatial rule the reward was always placed in either the east or west arm, while during the response rule the reward was placed in the light-on arm. Importantly, also during the spatial rule one of the two arms was lit, but did not necessarily indicate the location of reward. To prevent an odor-guided strategy pellet dust was scattered along the maze and pellet-filled cups invisible to the animal placed under both goal arms. On each recording day, the animal underwent behavioral blocks as follows: rest, rule 1 (remote recall), rest, pre-switch, switching, post-switch, rest, rule 2 (recent recall), rest. The analysis presented here was only performed on data recorded during the pre-switch, switching and post-switch blocks. Therefore only these blocks will be described in detail here. During the pre-switch block the animal had to collect reward based on the last rule of the previous day until reaching the performance criterion (see below). Then the rule was changed and reward had to be collected based on the new rule. The change in rule was not announced to the animal, which had to switch to the new rule through trial-and-error until performing to criterion. Trials performed after the rule change, but before the animal reached good performance comprised the switching block, while the post-switch block comprised all trials from the beginning of good performance (see *Methods-Behavior* for definition of good performance). The animal had to perform cross-modal switches, i.e., switches from spatial to light or light to spatial rule, never between the two spatial rules. While correct performance of a spatial rule involves two trajectories (e.g., go-east rule: north to east and south to east), correct performance of the light rule can involve any of the four trajectories. Therefore, the performance criterion for the spatial rule was set to 12/15 and for the light rule to 24/30 correct trials, ensuring similar number of light rule trials where the animal performed trajectories that matched those of the spatial rule.

Histology and Reconstruction of Recording Positions

After the final recording day tetrodes were not moved. Animals were administered ketamine/xylazine (7:3 ketamine (10%) and xylazine (2%), 0.1ml/100 g) and overdosed with pentobarbital (300mg/ml) before being transcardially perfused with 0.9% saline followed by 4% formaldehyde. Brains were extracted and stored in 4% formaldehyde. On the same day brains were transferred into 30% sucrose solution until sinking for cryoprotection. Finally, brains were quickly frozen, cut into coronal sections with a cryostat (50-60 μm), mounted on glass slides and stained with cresyl violet. The positions of tetrode tips were determined from stained sections and cells recorded from tetrodes outside mPFC were excluded from analysis. For cells recorded from HPC tetrodes the presence of SWRs in the field recordings served as inclusion criteria.

Data Acquisition

The extracellular electric signals from tetrodes were pre-amplified using a headstage (4 x 32 channels, Axona Ltd, St. Albans, Hertfordshire, UK). The amplified local field potential and multiple-unit activity were continuously digitized at 24 kHz using a 128-channel data acquisition system (Axona Ltd). Two red LED bundles mounted on the preamplifier head-stage were used to track the location of the animal. Every day before recording, HPC tetrodes were moved optimizing the yield of recorded cells. Additionally, mPFC tetrodes were lowered every day by ~30-50 μm to ensure recording of a new population of cells.

Spike Sorting and Unit Classification

Clustering of spikes and unit isolation procedures were described previously (Csicsvari et al., 1998). Briefly, the raw data was re-sampled to 20 kHz and the power in the 800-9000 Hz range was computed for sliding windows (12.8 ms). Action potentials with a power of > 5 standard deviations (SD) from the baseline mean were selected and their spike features extracted with principal components analysis. Action potentials were then grouped into multiple putative units based on their spike features using an automatic clustering software (http://klustakwik.sourceforge.net; Harris et al., 2000). The generated clusters were then manually refined using a graphical cluster-cutting program and only units with clear refractory periods in their autocorrelation, well-defined cluster boundaries and stability over time were used for further analysis. An isolation distance (based on Mahalanobis distance) was calculated to ensure that spike clusters did not overlap (Harris et al., 2000). Putative excitatory pyramidal cells and inhibitory interneurons were discriminated using their auto-correlograms, firing rates and waveforms. In our analysis we included 1025 HPC and 817 (of which 368 right and 449 left hemisphere) mPFC pyramidal cells.

QUANTIFICATION AND STATISTICAL ANALYSIS

Behavior

The analysis presented here was performed on data recorded during the pre-switch, switching and post-switch blocks. All trials before the rule change comprised the pre-switch block. Trials performed after the rule change, but before the animal reached good performance comprised the switching block. The beginning of good performance (bgp) was defined as the center index after rule change where the error rate over five consecutive trials dropped to zero.

$$bgp = argmin_{t \in \{trials > switch\}} \left\{ \sum_{k=-2}^{2} error(t+k) == 0 \right\}$$

Linearized Position

To linearize the behavior of the animal, we calculated the distance from the center from the 2D spatial position of the animal. This way a "V-shaped" positive function for each trial was obtained. For each position before the center (i.e., before the global minimum) we subtracted the minimum and then changed the sign. Then, 100 was added to every position to obtain a positive measure of the relative position of the animal between start (0 cm) and goal (200 cm). The center corresponded to 100 cm.

Linearized Firing Rate Maps

The linearized 1D maze was divided into 20 spatial bins of 10 cm. Using the linearized position, an occupancy map was generated by computing the amount of time the animal spent in each spatial bin during running periods (speed filtered > 5 cm/s). We then counted the number of spikes a cell emitted in each spatial bin (also speed filtered, > 5 cm/s) and divided that number by the occupancy time. We then smoothed these vectors with a Gaussian filter with a SD of 1 bin.

Population Vector Similarity

For each cell a rate map was computed separately from spiking data from each condition (e.g., south arm, correct trials,...). Rate maps for correct and error trials were computed from goal arm spiking data only. For each condition, rate maps of all cells in a given session were stacked along the z axis to obtain a set of population vectors, one for each spatial bin Then, for a population vector pair (i.e., one from each condition) for a given spatial bin the Pearson correlation was computed to quantify their similarity.

Bayesian Decoding

The position of the animal was decoded using a Bayesian decoding algorithm on the spiking vectors of cells and their expected firing rate given by their rate maps (Zhang et al., 1998). Spiking neurons were assumed to follow a Poissonian law and act independently of each other. The position of the animal in a given time window of length τ is denoted as x. The number of spikes of one particular cell is denoted as σ_i and the activity of all cells considered in the given time window as $\overline{\sigma} = \{\sigma_1, \sigma_2, \ldots\}$. For each cell i and position x, we obtain the measured firing rate λ_x^i from the rate maps. Using the assumptions and the introduced notation we have:

$$P\left(\overline{\sigma} \middle| x\right) = \prod_i P(\sigma_i \mid x) = \prod_i Poiss\left(\sigma_i \mid \lambda_x^i \tau\right)$$

And, using the Bayesian rule:

$$P\left(x \middle| \overline{\sigma}\right) = P\left(\overline{\sigma} \middle| x\right) \; P(x) \; / \; P(\sigma)$$

For our computations we used a flat prior (i.e., $P(x)$ was uniform). The term $P(\sigma)$ did not need to be computed because one can enforce $\sum_x P(x \mid \overline{\sigma}) = 1$. The decoded position then is:

$$BD(\overline{\sigma}) = argmax_x \; P\left(x \mid \overline{\sigma}\right) = argmax_x \prod_i Poiss\left(\sigma_i \mid \lambda_x^i \tau\right)$$

The decoded position was the position with the maximum likelihood. For a given window, the likelihoods for all positions were normalized by dividing each by the maximum likelihood of that window. However, in cases where we needed to consider the posterior as a properly normalized probability distribution, we normalized the likelihood by the sum. The same estimated firing probabilities of each cell in each bin were used for the pre-switch, switching and post-switch blocks.

Confusion Matrix

The confusion matrix was computed as follows: given a vector of real positions *rp* and a vector of decoded positions *dp* with same length N, both binned into bins of 10 cm size, we initialized a 20 × 20 matrix (CM) to zero where rows corresponded to the real position and the columns corresponded to the decoded position and filled it:

$$CM[i,j] = \sum_{k=0}^{N} (rp_k == i) \; and \; (dp_k == j)$$

Each row was then normalized so the sum of its elements equals 1.

Cross-validation

The decoding quality of a neural ensemble was tested on rate maps constructed from data that was not used to train the model. For spatial coding analysis, only running periods were considered (speed > 5 cm/s) and two thirds of the trials were used to compute the rate maps and the remaining third was used to decode the position of the animal. For analysis on non-local replay, rate maps were constructed from running (speed > 5 cm/s) and immobility (speed < 5 cm/s) periods and using all trials except for the one on which positional decoding was performed. For trajectory replay analysis during immobility periods, rate maps were constructed from running periods. The decoding error was the absolute value of the difference between recorded and decoded position.

Decoding Windows

The spiking vector on which positional decoding is performed can be chosen in different ways. The standard approach is to fix the time lengths of spiking vectors. Another approach is to select spiking vectors containing a minimum amount of spikes, without fixing the time, or selecting a minimum number of cells spiking. We binned the spiking activity of each cell in time bins of 25.6 ms (i.e., sampling-window), aligned with the behavior tracking. The 25.6 ms spiking vectors were then summed up until the request of a certain time window length, number of spikes or number of cells participating was met. We used fixed time windows of 256 ms for the detection of non-local positions during immobility. For the detection of trajectory replay we used windows with fixed number of spikes (i.e., spiking-window). The number of spikes contained in a window was selected for every session separately as follows: we computed the average decoding error during running periods (using the cross-validation method as described above) and determined the minimum number of spikes that allowed us to decode the position of the animal with a decoding error of less than 30 cm.

Decoding of Trial Outcome and Goal Arm Identity

We performed Bayesian decoding to determine whether HPC and mPFC firing codes for correct and error trials. First, we computed rate maps of randomly selected 80% of the spiking data occurring in the goal arms, separately for correct and error trials. Then for the remaining 20% of goal arm spiking data we performed decoding with the correct and error rate maps. The rate map that resulted in the highest likelihood determined the decoding outcome. The decoding outcome was then compared with the actual trial outcome and the percentage of trials where the decoding matched the trial outcome was plotted. Bayesian decoding was also performed to determine whether HPC and mPFC firing codes for goal arm identity (i.e., east or west). The same procedure as above was applied, just that rate maps were constructed either from east or west spiking data.

Trajectory Score

To ensure that the replayed trajectories observed during immobility were not statistical noise of neurons, we computed a trajectory score that consisted of a decoding probability and smoothness score. Each immobility period was divided into spiking-windows and, using the Bayesian algorithm, we obtained a normalized likelihood function on the binned linearized positions for each spiking-window. Every four consecutive spiking-windows (i.e., the first three windows of a four spiking-window event will be the same as the last three windows of the previous event) that passed the continuity (no jumps > 6 bins) and speed (absolute speed > 20 cm/s) criteria were analyzed and forward and backward replaying trajectories were detected separately. The decoding probability score is a measure based on a previous study (Xu et al., 2019). Starting from the first window of a four-window event, we computed the probability of transitioning to the next with a velocity given by the slope of a line fitted on the four decoded positions. We multiplied the likelihood with a Gaussian kernel that moved by the corresponding amount of spatial bins to the left (backward) or to the right (forward). If the limit was reached (< 0 or > 20) the kernel was set to zero. We computed this measure for all windows of a four-window event and averaged the result to obtain the decoding probability score ranging between 0 (least optimal) and 1 (most optimal). The smoothness score measures how probable it is to find a permutation of the given four spiking-window event that has a distance between neighboring windows smaller than the original event. In the case of four or five windows, 24 or 120 permutations are possible. The number of permutations that had an average distance between neighboring events smaller than the real event, was divided by the total number of permutations, yielding a number from 0 (most optimal) to 1 (least optimal). For events longer than five windows, since it becomes computationally prohibitive to compute all the permutations, we randomly sampled 100 permutations and computed the score in the same way on those 100 samples. The smoothness score was then subtracted from the decoding probability score to obtain our final trajectory score. This way we penalized events that showed non-sequential activity.

Shuffling Procedure

The trajectory score of each four (or longer) spiking-window event was compared to its own shuffled distribution. The shuffled distribution was generated by decoding the same event (as described above), but this time using 200 shuffled rate maps, and computing 200 trajectory scores. We used a place field rotation shuffling. Since we are working with 1D rate maps, this entails drawing a random integer between 3 and 17, and wrapping the map of each single neuron (independently) by that number of bins around the origin. For some analyses we additionally employed a cell identity shuffling procedure, where the cell labeling of rate maps was permutated but otherwise were kept intact.

Trajectory Selection and Concatenation

Those four spiking-window events that had a unified trajectory score above the 95th percentile of their own shuffled distribution were selected. A given significant four spiking-window event was extended by the consecutive four spiking-window event (which overlaps by three windows) if the consecutive event also had a significant trajectory score, the speed and jump size matched our criteria and the newly calculated trajectory score still passed the 95th percentile of its own shuffled distribution. This way we extended a four spiking-window to a five spiking-window event and the same procedure was then performed on the extended event.

Cross-correlation

Cross-correlation on HPC and mPFC trajectory replay events was performed to investigate their temporal alignment. For each brain area we considered a time series, aligned with behavior (25.6 ms), filled with zeros in case of immobility, ones for significant trajectory

replay, and NaNs otherwise. The mPFC time series was incrementally shifted to the left or right relative to the HPC time series until a maximum of 20 bins. For each shift, the Pearson's correlation between the two time series was computed excluding the NaNs (and thus mobility and non-task periods). To compare the cross-correlation to shuffled data, trajectory events were randomly shifted forward or backward within their immobility period 100 times, and each time a cross-correlation computed.

Additional Quantification of Linear Replay Quality

Additionally to the trajectory score, we performed further analyses to quantify linear replay with methods employed in previous replay studies (Davidson et al., 2009; Gupta et al., 2010). We will briefly outline the main idea and fundamental steps of the algorithms and then refer the reader to the original publications for further details. The algorithm in the Gupta et al. study consists of assigning one discrete position to each cell, which corresponds to the maze position with the highest firing rate observed during running. For each candidate replay event the spiking of cells is divided into 25.6 ms time windows. Then, for each spike in each window, one determines if the spikes in the subsequent three time windows correspond to maze positions that precede or succeed the position encoded by the current spike and then averages the outcome. Therefore, this score measures the propensity of sequential firing of cells in consecutive short time windows, separately for forward and reverse replay.

To compute the score as in the Davidson et al. study, for each candidate replay event the posterior probability over linear bins was computed for each time window. These are stacked to form a T x N matrix, where T is the number of time windows and N the number of spatial bins. Then, all possible linear trajectories spanning the entire time by space matrix with a minimum speed are considered. Here, were considered speeds in the range of 5 to 40 cm/step, both forward and reverse. For each of the above-mentioned trajectories, all those probabilities were summed that were crossed by the selected trajectory or were less than 10 cm away, for a total of 20 cm. Only the sum of probabilities of the best scoring trajectory was returned.

Decoding of Goal Arms within Replay Events

Although replay events were detected with 1D rate maps, we nonetheless tried to determine whether a trajectory replay event recapitulated the current or previous trial's goal arm. The 2D position was decoded for every spiking-window of a replay event. If at least two windows encoded a goal arm, the goal arm identity that was being replayed was determined by a majority rule.

Population Firing Rate during Replay

The population firing was binned into 25.6 ms windows and Z-scored relative to the firing in the entire session. Then for every replay event, the beginning of the spiking-window with the highest firing rate was determined and the population firing rate vector extracted within an interval of −460.8 ms and +460.8 ms (i.e., 18 bins before and after the spiking-window with highest firing rate). The mean and standard error of the mean (SEM) over all extracted firing rate vectors was plotted.

Oscillation Analyses

The local field potential signal (5 kHz) of one channel per HPC tetrode was subtracted with the signal of a reference channel that did not contain SWRs. For steps of 20 ms, 240 ms chunks of the signal were Fast-Fourier transformed and band-pass filtered between 150 and 250 Hz. The root mean square was then computed to obtain the ripple-band power for that channel. The mean ripple-band power was obtained by averaging over all channels. SWR events were detected by detecting periods where the ripple-band power was above 5 SD of the mean and the start and end times were also determined (1.5 SD). The same procedure was done for cortical ripples, however right hemisphere mPFC channels were subtracted with the signal of a left hemisphere channel and the signal was filtered between 100 and 150 Hz (Khodagholy et al., 2017). To obtain the gamma power the signal was filtered between 30 and 90 Hz. No reference channel was subtracted from the signal and for mPFC gamma power only right hemisphere channels were considered. For every replay event an immobility period of the same length was randomly selected and the respective type of network activity computed over these periods for comparison.

Behavioral Correlates

The immobility time in a trial or experimental block can vary depending on how long it takes for the animal to reach and consume the reward or to make a decision. We therefore computed a trajectory rate for every trial (i.e., number of trajectory replay events divided by the immobility time) instead of counting the number of trajectories per trial. We computed a trajectory rate separately for the center and goal area (i.e., at a linearized position of < 110 cm or > 110cm, respectively). For the center, trajectory replay events were also divided into ahead and behind replaying trajectories. Depending on the analysis, an average rate was computed for every session from trials within the switching or within the switching and the post-switch block. When correlating the trajectory rate with the number of trials needed to switch to the new rule, we employed a logarithmic transformation of the trajectory rate. Insets in correlation figures show boxplots of 100 r values computed from 100 correlations of 80% randomly sampled data points of the data. The number of immobility periods was computed by counting the number of periods the animal had a speed of lower than 5 cm/s and dividing that number by the total number of trials. The rate of SWR events was computed for every trial by counting the number of SWR events during immobility periods and normalizing by the total time spent immobile during that trial. An average SWR event rate per trial was then computed.

Statistical Analysis

Analyses were performed using custom written Python scripts (https://www.python.org). Non-parametric and two-tailed tests were used throughout unless stated otherwise. For paired comparisons Wilcoxon signed-rank test was used, while Mann-Whitney *U* test was used for unpaired data. For the analysis on trajectory replay rate in error and correct trials, the Wilcoxon signed-rank test was used with the Bonferroni-Holm correction for multiple comparisons. For other multiple group testing the Kruskal-Wallis test was used. The two-sample Kolmogorov-Smirnov test was used to compare distributions. A one-sample t test tested for changes in firing rates. Correlations were performed by calculating the linear least-squares regression and the significance was determined using Spearman's rank-order correlation. Correlations determining population vector similarities and cross-correlations were performed with Pearson's correlation. All statistical tests are reported in the text and appropriate figure legends ($^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$, $^{****}p < 0.00001$). In bar plots the mean and SEM are shown.

DATA AND CODE AVAILABILITY

Data and software used in this study will be made available upon request by contacting the lead contact, Jozsef Csicsvari (jozsef.csicsvari@ist.ac.at).