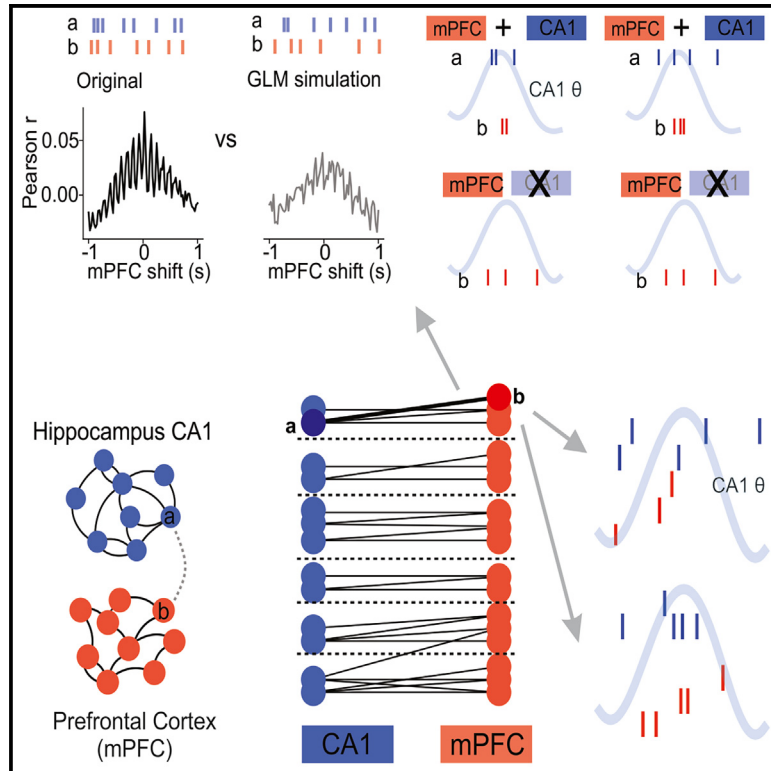


Theta oscillations as a substrate for medial prefrontal-hippocampal assembly interactions

Graphical abstract



Authors

Michele Nardin, Karola Kaefer, Federico Stella, Jozsef Csicsvari

Correspondence

nardinm@janelia.hhmi.org (M.N.), jozsef.csicsvari@ist.ac.at (J.C.)

In brief

Nardin et al. identify functionally coupled hippocampal CA1-medial prefrontal cortex (mPFC) assemblies. During the synchronized activity of an assembly, mPFC cells enhance their phase locking to CA1 theta oscillations and fire along specific theta sequences. Thus, theta sequential firing of mPFC cells can signal transient interactions with CA1 assembly partners.

Highlights

- mPFC spatial coding lags behind the hippocampus with a behavior-dependent delay
- CA1-mPFC cells synchronize independently from common place coding or oscillations
- Enhanced mPFC theta phase locking signals interaction with specific CA1 assemblies
- Interacting CA1-mPFC assemblies can exhibit independent theta sequences



Article

Theta oscillations as a substrate for medial prefrontal-hippocampal assembly interactions

Michele Nardin,^{1,2,*} Karola Kaefer,^{1,3} Federico Stella,^{1,3} and Jozsef Csicsvari^{1,4,*}¹IST Austria, 3400 Klosterneuburg, Austria²Janelia Research Campus, Howard Hughes Medical Institute, Ashburn, VA 20147, USA³Present address: Donders Institute, Radboud University, 6500GL Nijmegen, the Netherlands⁴Lead contact*Correspondence: nardinm@janelia.hhmi.org (M.N.), jozsef.csicsvari@ist.ac.at (J.C.)<https://doi.org/10.1016/j.celrep.2023.113015>**SUMMARY**

The execution of cognitive functions requires coordinated circuit activity across different brain areas that involves the associated firing of neuronal assemblies. Here, we tested the circuit mechanism behind assembly interactions between the hippocampus and the medial prefrontal cortex (mPFC) of adult rats by recording neuronal populations during a rule-switching task. We identified functionally coupled CA1-mPFC cells that synchronized their activity beyond that expected from common spatial coding or oscillatory firing. When such cell pairs fired together, the mPFC cell strongly phase locked to CA1 theta oscillations and maintained consistent theta firing phases, independent of the theta timing of their CA1 counterpart. These functionally connected CA1-mPFC cells formed interconnected assemblies. While firing together with their CA1 assembly partners, mPFC cells fired along specific theta sequences. Our results suggest that upregulated theta oscillatory firing of mPFC cells can signal transient interactions with specific CA1 assemblies, thus enabling distributed computations.

INTRODUCTION

Hebb's postulate¹ that coordinated activity enables neurons to associate with each other and ultimately establish cell assemblies influenced neuroscience from the cellular to systems levels for decades. Recent developments in multi-neuronal recordings have provided the means to test the cell assembly hypothesis over larger neuronal populations.² While neurons can exhibit highly correlated activity, similar firing responses can also arise from common sensory inputs. The reactivation of population firing patterns during sleep offered a potential testing ground for assemblies, considering the diminished sensory input to cortical areas in sleep.³ It is also possible to investigate assemblies outside sleep by applying statistical methods such as noise correlations that were developed initially in sensory neuroscience to take into account the common sensory input.^{4–6} However, one may need to consider additional non-sensory factors, such as network oscillations, that can synchronize large populations of cells.⁷

To understand the contribution of cell assemblies to cognitive functions at the systems level, one needs to determine how assemblies in different brain areas work together and how the activity of assemblies propagates across regions. The synfire chain theory provided a potential mechanism in which the synchronized activity of associated assemblies is part of a synfire chain that propagates across multiple brain regions.⁸ As for within-region assemblies, sleep reactivation provides a means to identify related cross-regional assemblies independent of a common

sensory drive. Reactivation in the hippocampus is linked to sharp wave-ripple patterns (SWRs),⁹ and synchronized network patterns during SWRs induce large-scale, brain-wide synchronization.^{10,11} Furthermore, reactivated spatial assembly activity in one region can recruit assemblies with similar coding properties in another region.^{12–14} However, the reactivation of place-coding assemblies is often not coordinated across regions, or the assemblies recruited in a downstream area in sleep do not correspond with waking assemblies. While transient sequential activation of spatially firing cells in multiple brain regions can express trajectories, these sequential reactivation patterns detected in different regions are not coordinated,^{15,16} although, in certain conditions, coherent sequences can be observed.¹⁷ Furthermore, neurons downstream may decode non-local spatial information derived from hippocampal reactivated patterns in sleep but not during waking behavior. In agreement with this, some medial prefrontal cortex (mPFC) cells increase their firing when trajectories of a specific maze arm are reactivated in the hippocampus without these cells showing increased firing at that arm during behavior.¹⁸

It is expected that assemblies exhibit coordinated activity across regions during behavior as well, when their activity is also co-modulated by a common sensory drive. In this condition, as is the case of within-region assemblies, one needs to take into account common factors that can lead to coherent activity across regions without one area influencing the other. Here, we tested whether assemblies coordinate their activity across the hippocampus and the mPFC while both regions engaged in



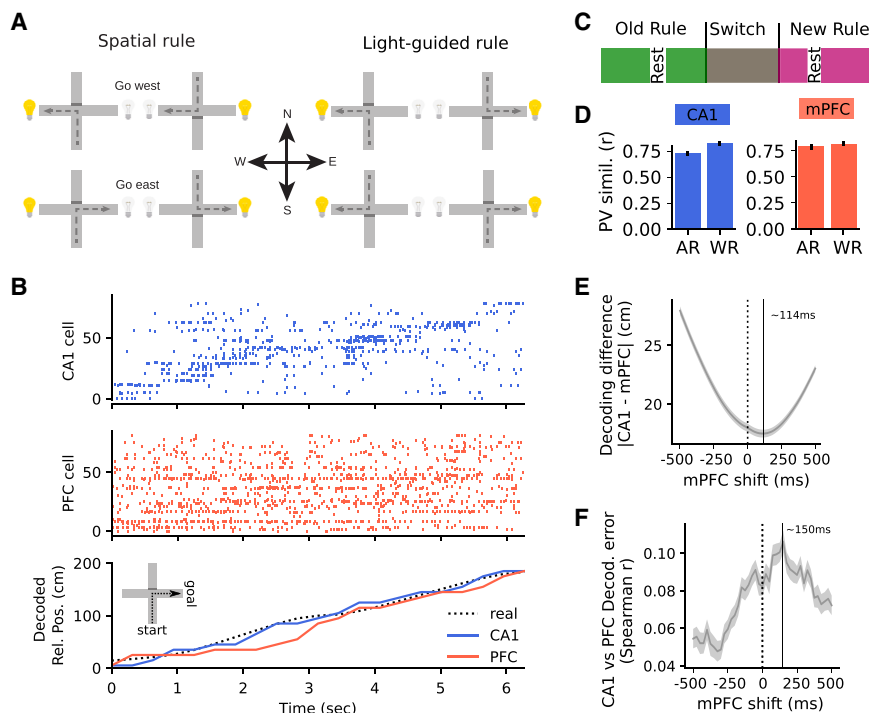


Figure 1. Delayed spatial coding in the mPFC relative to the hippocampus

(A) Behavioral rules applied during the task. (B) Population firing during a single trial. Top two panels: raster plot of CA1 and mPFC population firing. Each line represents the firing of a separate cell. Bottom panel: position prediction from population activity and the actual position of the animal. (C) Experimental paradigm: three sessions were separated by rest periods, and the rule is switched in the middle session. (D) Similarity of average firing rate population vectors (PVs) across rules (ARs) vs. within rules (WRs) for CA1 ($n = 530$) and mPFC ($n = 477$) putative principal cells. Only CA1 PV similarity is different ARs vs. WRs (Tukey's HSD post hoc AR vs. WR, CA1 $p = 0.019$; mPFC, $p = 0.95$, all other $p > 0.08$). Error bars: \pm SEM. (E and F) Correlation of predicted locations and decoding errors as a function of mPFC time shift. Shaded area: 95% confidence interval (CI), solid lines: peak, $n = 13$ sessions in four animals.

spatial coding. We were able to identify CA1-mPFC assemblies that synchronize their activity independently from common sensory or oscillatory inputs. We also showed that when these assemblies were active, the theta-oscillation-related temporal firing patterns of mPFC cells signaled interactions with their CA1 partners.

RESULTS

We analyzed simultaneously recorded neuronal population activity data in the dorsal hippocampal CA1 area and the prelimbic mPFC during a task requiring a switch between spatial and cue-guided rules on a plus maze¹⁵ (Figures 1A and 1B). Inactivation studies suggest that CA1-mPFC spatial representations mutually influence each other,¹⁹ even though they can code for spatial information in different forms.^{15,20,21} In our experiments, mPFC populations encoded relative maze distance more accurately than actual 2D positions.¹⁵ In a plus maze, such an arm-independent distance code requires similar spatial firing on opposite arms. In agreement with this, we found that a large proportion of mPFC cells had symmetric firing fields (Figure S1B). Although the mPFC is a crucial region for rule switching,²² spatial representations showed only marginal differences across the two rules and exhibited significant differences in the hippocampus only (Figure 1D), and even in error trials were representations similar (Figure S1E). To quantify the encoding dynamics of spatial representations, we tested whether spatial coding in one region lagged behind the other. Moreover, to infer whether spatial information in one region is used by the other, we assessed the correlation of decoding errors.²³ The best alignment of decoded positions was seen when the mPFC lagged behind

that assesses information exchange without assuming spatial coding also indicated a preference for CA1-to-PFC interactions (Figure S1C). These findings show that CA1 spatial coding emerges earlier and that the more abstract distance code in the mPFC may use information derived from CA1 cognitive maps.

To test for the presence of interregional assembly interactions, we first examined whether functional interactions can be identified at the cellular level. Neurons in both regions code for place while also engaging in coherent oscillations.^{23,25} Therefore, we calculated functional cross-correlations of CA1-mPFC cell pairs that measured the strength of correlated firing that occurred independently of spatial firing, oscillations, or other possible common drivers. For each cell, we fitted a generalized linear model (GLM)²⁶ that included spatial position, trajectory, speed, theta-oscillatory firing, spiking history, and the within-area spiking of other cells (Figure S2; see STAR Methods). The cross-correlogram calculated from the data was compared (Z scored) with the distribution of cross-correlograms computed on simulated spiking (10,000 times) using the GML model (Figures 2A and 2B). In testing different temporal offsets, we found that the maximal percentage of cell pairs with strong functional correlation ($>4.5 Z$, $p < 0.00001$) occurred with a ≈ 50 ms mPFC delay relative to the CA1, and such strong functional correlations occurred mostly within a 125 ms time window (Figure 2C). We therefore selected functionally coupled (FnC) cell pairs that had $>4.5 Z$ peak functional cross-correlation within [0,125] ms (Figures 2B, S3A, and S3B).

Although FnC cells synchronized their firing independently of a common spatial drive, CA1 firing fields tended to be on the same arms as those of their FnC mPFC counterparts (Figures 2B, S4,

the CA1 by ≈ 114 ms (Figures 1E and S1A), and decoding errors also correlated the strongest at ≈ 150 ms mPFC delay (Figure 1F). A transfer entropy measure²⁴

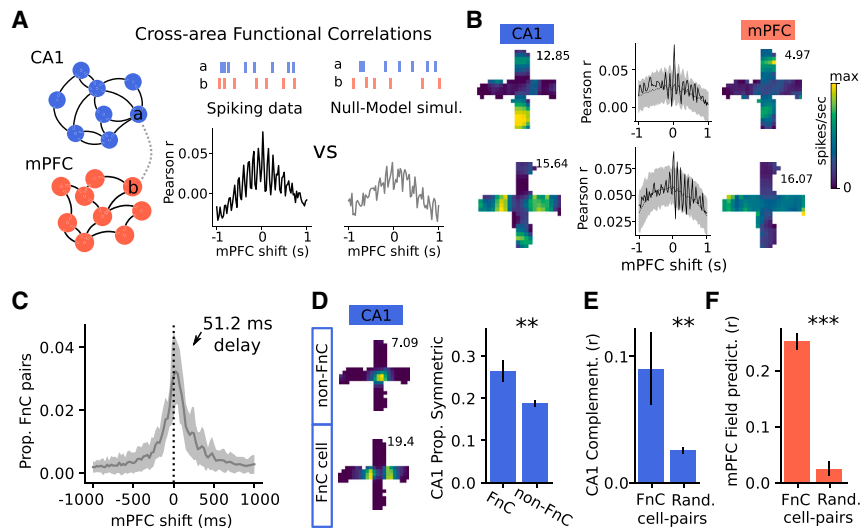


Figure 2. CA1-mPFC functional correlations
(A) Functional correlation (FNC) calculations: real data cross-correlations were compared to a null distribution from 10,000 GLM simulated cross-correlations.
(B) Example cross-correlations with significant interaction ($Z > 4.5$ relative to GLM). Left and right: CA1 and mPFC firing rate map (inset number: max firing rate). Center: cell-pair cross-correlograms, real data (solid black line); null model mean (dotted back) ± 4.5 SD (shaded area).
(C) Proportion of significant Fnc pairs ($Z > 4.5$) as a function of mPFC delay. Shaded area: 95% CI, $n = 23,482$ pairs across 13 sessions.
(D) Left: example place fields of non-FnC and FnC cells. Right: proportion of symmetric place fields for FnC and non-FnC cells.
(E) Complementariness (i.e., a tendency to fire at opposite arms, see STAR Methods) of rate maps of the two CA1 units with the strongest FNC to the same mPFC unit vs. random.
(F) Prediction of mPFC firing rate map by Fnc CA1 firing rate maps. mPFC map was predicted by the FNC weighted sum of the CA1 maps (see STAR Methods) vs. random.
(D–F) Mann-Whitney U test: ** $p < 0.01$, *** $p < 0.001$. Error bars: $\pm 95\%$ CI.

and S5). Furthermore, CA1 Fnc cells either exhibited a tendency for symmetrical place fields (Figures 2D and S1D) or pairs of them fired on the opposite arms (in a complementary fashion) (Figures 2E and S4G). Moreover, the CA1 place fields predicted the firing field of the Fnc mPFC cell (see STAR Methods; Figures 2F and S4H), and opposite mPFC prediction of CA1 place fields was also possible (Figure S4I). A computational model of spatial information transfer that maximizes mutual information (see STAR Methods) predicted these connectivity rules (Figures S4A–S4F, S4J, and S4K).

The cross-correlations among Fnc cell pairs exhibited theta-modulated peaks (Figures 2B and S5), indicating the mPFC cells phase locked to hippocampal theta oscillations. Previously it was shown that mPFC theta-oscillatory coupling to the CA1 changes in strength depending on the cognitive demand.^{25,27} Therefore, we examined theta phase locking of mPFC cells separately at different task phases in the start and goal arms. Fnc mPFC cells exhibited stronger phase locking to CA1 theta oscillations than non-Fnc cells both on the start and the goal arms, while CA1 pyramidal cells did not show such differences (two-way ANOVA, start arm: $F(\text{Fnc}/\text{non-Fnc}) = 6.85$, $df = 1$, $p < 0.01$, goal arm: $F(\text{Fnc}/\text{non-Fnc}) = 4.71$, $df = 1$, $p < 0.03$; Holm post-hoc t test, Fnc vs. non-Fnc, start arm: CA1 $p = 1$, mPFC $p = 0.00124$; goal arm: CA1: $p = 1$, mPFC: $p = 0.0302$; Figures 3A and 3B). The phase locking of mPFC cells improves when it is referenced to a time-delayed theta-oscillatory local field potential (LFP) signal from the CA1.²⁸ Therefore, we calculated the optimal time delay for the theta phase locking of mPFC cells. The median delay was 23 ms in the start arm and 47 ms in the goal arm (Figure 3C). Considering the delay differences we saw for theta phase locking, we also recalculated the proportion of significant functional correlations separately for start and goal arms, which showed delays of 25 (start arm) vs. 75 ms (goal arm)

(Figure 3D). Although theta phase locking predicted whether a mPFC cell functionally connected to a CA1 cell, mPFC cell recruitment to hippocampal SWRs and gamma oscillations was similar when Fnc and non-Fnc cells were compared (Kolmogorov-Smirnov [KS] test, all $p > 0.435$; Figure S6A).

Next, we examined the underlying mechanisms that allow for cross-area interactions. Theta-modulated cross-correlations can be seen when one or both cells phase lock to theta oscillations but otherwise fire independently. However, in our Fnc cells, the theta-modulated firing of cells was accounted for by the GLM (Figures 2A and 4A); hence, we do not expect theta modulation on the functional correlations of independently firing theta-modulated cells. Therefore, theta-modulated functional correlations could indicate a transient increase in the theta phase locking of Fnc cells when they synchronize their activity. This enhanced phase locking of the mPFC cell can signal firing epochs, in which it synchronizes its firing with its CA1 Fnc counterpart. To directly test this, we examined whether mPFC cells exhibited better phase locking to theta oscillations when they fired within a theta-cycle delay of their Fnc CA1 counterparts (Figures 4B and 4C). To quantify this effect, we also identified unpaired mPFC spikes, in which cases the mPFC cell fired independently from its CA1 counterpart (Figure 4C, bottom panels). Then, we randomly subsampled the unpaired mPFC spikes so that we got an equal amount to that of the paired spikes. In a bootstrapping procedure, we repeated the subsampling 200 times to estimate the mean and SD of the theta phase-locking strength of unpaired spikes and then Z scored the phase-locking strength of the paired spikes (Figure 4D). These Z scores of different mPFC cells showing the normalized phase-locking strength difference of paired and unpaired spikes exhibited a significant positive bias (Figure 4E, top row, binomial test, $p < 0.001$). Therefore, when Fnc cell pairs fired in the same theta

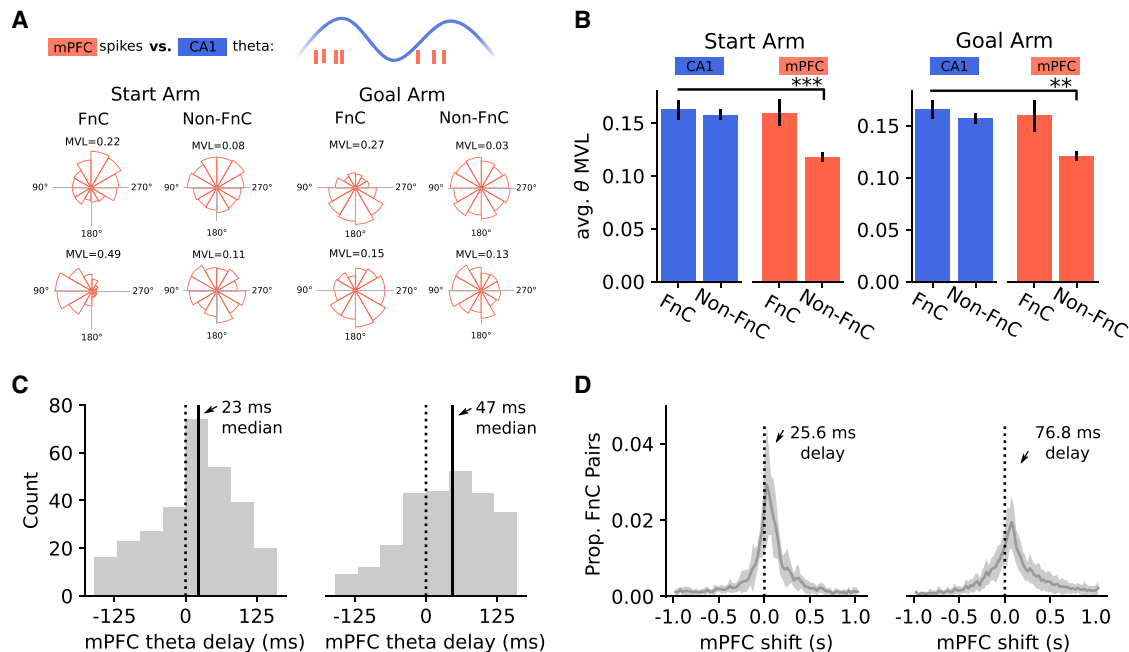


Figure 3. Coupling is delayed in the goal arm relative to the start arm

(A) Examples of theta spike phase histograms of FnC and non-FnC mPFC cells. MVL, mean theta vector length of the plotted cell quantifying place-locking strength.

(B) Average theta phase-locking strength of CA1 (blue) or mPFC (red) cells, comparing FnC vs. non-FnC (Holm post hoc t test, start arm: CA1 $p = 1$, mPFC $p = 0.00124$; goal arm: CA1: $p = 1$, mPFC: $p = 0.0302$).

(C) Optimal delay that produced the best theta locking of each mPFC cell, found by computing a phase histogram for each delay and measuring the mean vector length.²⁸ Mann-Whitney U test for start vs. goal arm optimal theta delay, $N = 477$, $p < 0.001$. KS test for the distributions, $p < 0.001$. Solid vertical line represents the median.

(D) Proportion of significant cell pairs ($Z > 4.5$) in start arm (left) vs. goal arm (right) as a function of mPFC delay. Shaded area represents the 99% CI for the mean. $n = 23,482$ pairs, $N = 13$ sessions.

cycle, mPFC theta phase locking was stronger than that during independent firing. Such a difference was significantly smaller when non-FnC pairs were tested (Figure 4E, bottom row, Mann Whitney U test, all $p < 0.0003$). For cell pairs detected with a lower threshold ($3.5 < Z < 4.5$), we still observed a stronger mPFC cell phase locking for CA1-paired firing compared with unpaired ones (Figure S3C).

Although we observed enhanced phase locking of mPFC cells when synchronized with their FnC CA1 counterpart, this does not imply that it was required to detect a functional coupling. Therefore, we verified that the enhanced precision of mPFC spike timing (due to strengthened theta phase locking) was a critical factor for observing CA1-mPFC functional coupling. We randomized the theta spike timing of the mPFC cells during synchronized firing with the FnC CA1 counterparts. The GLM-corrected cross-correlation peak (that accounted for theta-modulated cell firing) was reduced by $\approx 50\%$, and the number of significant FnC cells dropped to a quarter of the original number (Figures S7A–S7C). We also performed an additional spike-jittering procedure, where spikes of each mPFC neuron, within every theta cycle, are redrawn according to their theta-phase histogram (Figure S7D). This procedure ensured that mPFC spikes remained in the same theta cycle and also preserved every aspect of mPFC activity apart from fine, subtheta cycle

synchronization. Even in this conservative procedure, the number of detected significant CA1-mPFC FnC pairs drastically reduced (Figure S7E). These results demonstrate that the transient strengthening of mPFC phase locking to hippocampal theta can be a critical mechanism for CA1-mPFC functional coupling.

Although CA1 cells exhibit, on average, a phase preference for theta oscillations, they can also systematically change their theta phases during theta phase precession.²⁹ Considering that mPFC cells transiently phase locked to theta oscillations during synchronized firing with FnC CA1 counterparts, they may either maintain a consistent theta phase or mirror theta phase precession of their CA1 counterparts. We calculated CA1-mPFC functional cross-correlations separately for cases when the CA1 cell fired at different theta phases (Figure 4F). The plot of significant interactions now exhibited theta-modulated peaks, and the peaks shifted in the opposite direction relative to the CA1 cells' theta-phase change (Figure 4F). This indicates that the correlation delay compensated for the phase shift of the CA1 cell. To verify this, we calculated the average phase difference of FnC mPFC-CA1 pairs when CA1 cells fired at different theta phases. This showed a negative relationship: the forward movement of CA1 phases was compensated by a backward change of mPFC-CA1 phase differences, indicating that mPFC cells

maintained a relatively fixed theta phase (Figure 4G). These findings show that during periods of synchronized firing, each mPFC cell maintained a consistent theta phase and did not follow the theta-phase changes of their CA1 counterparts.

Gamma oscillations can synchronize distributed sensory representations³⁰ or spatial representations during memory encoding.³¹ Therefore, we tested whether transient synchronization of gamma (or other) oscillations contributed to the functional coupling. We compared the power spectrum and coherence of the CA1 and mPFC LFPs at times when FnC CA1-mPFC cells fired together or independently (Figure 5). No differences were seen in any of the frequency bands, suggesting that selective mPFC cell theta phase locking is the primary mechanism behind the CA1-mPFC functional coupling and that these coupling events do not influence the global oscillatory state of the circuits per se.

So far, we have identified functional correlations between individual CA1-mPFC pairs. Can this pairwise connectivity reveal organized interregional assemblies? We applied a bipartite spectral graph-clustering algorithm^{32,33} that divided cells into subgraphs. When we plotted the graph of FnC cells ordered by the graph-clustering algorithm, the different subgraphs could be clearly separated (Figure 6A), and the probability of a cell projecting out of a cluster was about a quarter of the probability of projecting within a cluster, i.e., cluster cleanliness ≈ 0.75 (Figure S3D, right panel). We examined whether these subgraphs (we refer to them as graph assemblies) met the criteria of functional assemblies. First, we confirmed that mPFC cells within a graph assembly exhibited a stronger functional correlation with each other than with mPFC cells of different assemblies (Figures 6D and 6E). We then identified times when graph assemblies were active by selecting theta cycles in which at least two mPFC and two CA1 cells from the same graph assembly fired together. When such a graph assembly was active, we found that mPFC theta phase locking was stronger as compared with out-of-assembly activity (Figures 6F and 6G, right). Finally, we found that assembly activity encoded spatial locations consistently over time and that the firing fields of different graph assemblies were different (Figures 6B and 6C).

We showed above that FnC mPFC cells tend to fire at relatively fixed theta phases during graph-assembly activation (Figures 4E–4G). In doing so, they all may fire at similar theta phases, or if they fire at different theta phases, they can organize their firing into sequential patterns. To test this, first we calculated the preferred theta phase of mPFC cells during assembly activation. These preferred theta phases spanned the entire theta cycle (Figures 6F and 6G, left). During graph-assembly activation, mPFC cell members may therefore exhibit a consistent temporal bias in their spike timing to theta oscillations. For each mPFC cell pair from the same graph assembly, we tested whether they exhibited a significant theta temporal bias during assembly activa-

tion, i.e., one cell consistently fired at an earlier theta phase than the other. Indeed, in total, 45% (start arm) and 42% (goal arm) of FnC mPFC pairs exhibited a significant bias (Figure 6H, binomial tests: all $p < 0.001$), confirming that, at the population level, a significant proportion of cells showed a theta-phase bias. Therefore, during assembly activation, a significant portion of mPFC cell pairs show a preferred temporal order in their activity, leading to sequentially biased firing patterns (Figure 6I). We expect such sequences to be independent of CA1 firing patterns, considering that mPFC cells decoupled from the theta spike timing of CA1 cells (Figures 4E–4G). Theta sequences in the hippocampus reflect the compressed time order at which place cells fire on a linear maze.^{34,35} If mPFC cells of a graph assembly exhibit similar maze-related activation sequences, the temporal activation order of mPFC cells (seen on their long-range cross-correlograms) should predict their theta-phase bias during graph-assembly activation. However, this was not the case (Figure 6J, top), suggesting that during assembly activation, mPFC cells can exhibit unique sequences that are independent of their longer time scale activation patterns driven by maze trajectories. As a verification, we also checked that mPFC cells in a graph assembly do not exhibit maze-trajectory-related theta sequences by plotting the relationship place field distance vs. theta-range cross-correlation delay,³⁴ which also does not show a relationship (Figure 6J, bottom start arm: Spearman $r = -0.09$, $p = 0.302$; goal arm: $r = 0.093$, $p = 0.389$). Hence, these sequences could represent unique non-spatial information and signal functional interactions with specific CA1 assemblies.

Finally, we examined the synchronization of graph assemblies at different task stages. We measured the average synchronization level (i.e., proportion active) of these start and goal assemblies at different relative maze locations. When correct and error trials were compared, a larger proportion of assemblies were active in error trials but only on the goal arms excluding the central stem (Figures 7A and 7B, >20 cm, three-way ANOVA $F(\text{position}) = 285.115798$, $df = 2$, $p < 0.0001$, $F(\text{error}) = 132.580$, $df = 1$, $p < 0.0001$, Tukey's honest significant difference [HSD] post hoc, all $p < 0.04$). Next, we compared assembly synchronization in three subsequent trials during rule switching: the trial when the animal followed the new rule for the first time and the trials right before and right after that (Figure 7C). In this case, we saw increased assembly synchronization of start-arm assemblies before the strategy switch (compared with the other two trials) at the end of the arm (Figure 7D, three-way ANOVA $F(\text{position}) = 19.627883$, $df = 19$, $p < 0.0001$; $F(\text{switch}) = 28.887530$, $df = 2$, $p < 0.0001$; Tukey's HSD post hoc, all $p < 0.02$), while for goal-arm assemblies, only the before-switch and strategy switch trials were significantly different (Figure 7E, Tukey's HSD post hoc, $p = 0.0236$). This finding is probably related to the error trial results (Figure 7B), considering that the before-switch trial was often an error trial. Finally, in the center

(E) Top: Z scored MVL of mPFC cells during paired firing relative to the subsampled unpaired firing (see D). Bottom: same as in the top but for randomly selected non-FnC pairs. The Z scored MVLs were significantly larger for FnC pairs than for non-FnC ones (***, Mann Whitney U test, all $p < 0.0003$, start: FnC pairs = 444, non FnC = 693; goal: FnC pairs = 397, non-FnC pairs = 632).

(F) Left: schema of theta-dependent FNC calculation. Arrows point to cross-correlograms computed at different theta phases. Right: proportion of $z > 3$ FnC pairs as a function of CA1 theta phase and mPFC time shift.

(G) Average mPFC-CA1 theta-phase difference at different CA1 firing phases. Shaded area: 95% CI.

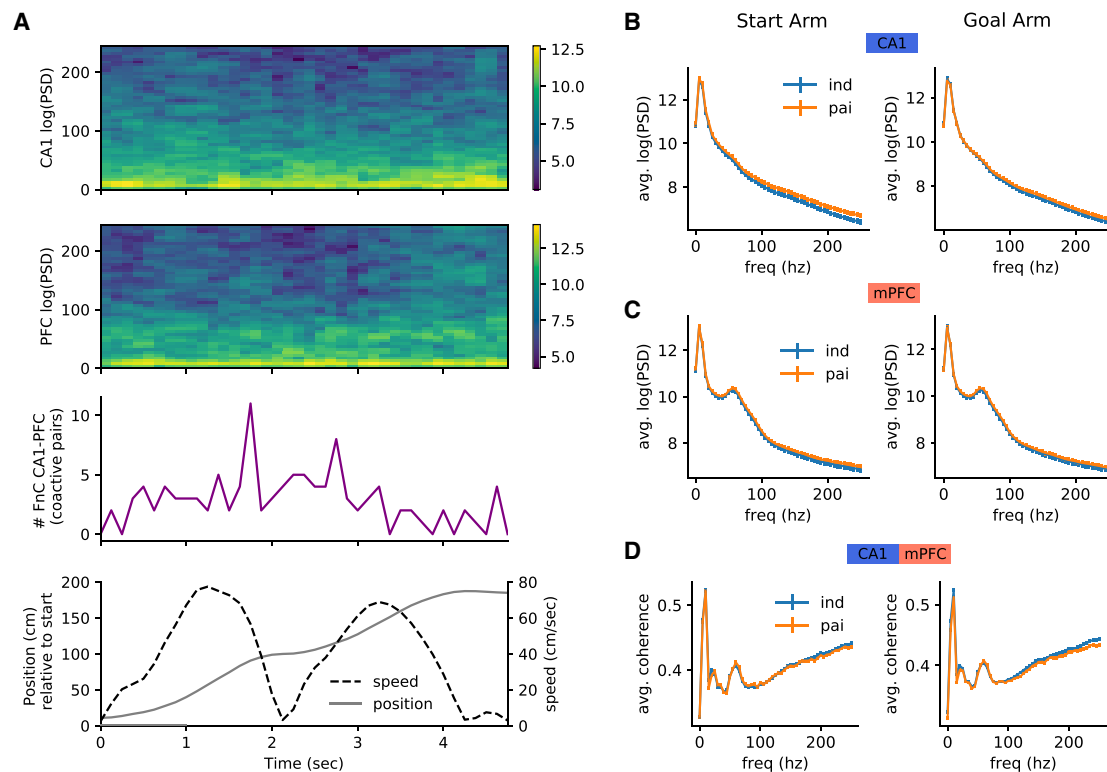


Figure 5. Paired firing of mPFC-CA1 FnC pairs is not accompanied by significant changes in underlying oscillations

(A) Illustration of the power spectra variation of the LFP during a single trial. In addition, the figure also plots the proportion of active FnC cells and the position and speed profile of the animal.

(B–D) Average power spectrum density (PSD) function (B and C) and coherence (D) for paired (orange) vs. independent (blue) activity of FnC cell pairs. PSDs were computed using Welch’s method on 400 ms time segments centered on mid-theta cycles using a Hanning window and constant detrending. The power and coherence were not significantly different in the paired condition compared with the independent conditions for theta (5–10 Hz), gamma (40–80 Hz), and ripple (150–250 Hz) frequency bands (power: CA1 all $p > 0.41$; mPFC all $p > 0.48$; coherence all $p > 0.69$). Error bars: \pm SEM.

of the maze, start-arm assemblies synchronized more strongly during the switch than before (Figure 7C, Tukey’s HSD post hoc, $p = 0.001$), while between the switch and after the switch, there was only a trend for significance (Tukey’s HSD post hoc, $p = 0.0575$). These results showed that graph-assembly synchronization could vary at certain behavioral conditions and could be modulated by errors or rule switching. Periods of strongly synchronized graph-assembly activation may correspond to a special network state marked by the enhancement of certain network oscillations. However, hippocampal and mPFC LFP power spectra did not show differences at periods of low- and high-synchrony assembly activation when the bottom and top 20% assembly synchrony levels were compared (Figures S6B and S6C, Mann-Whitney U test, CA1 all $p > 0.23$, mPFC all $p > 0.11$).

DISCUSSION

Here, we identified a subgroup of mPFC and CA1 cells that synchronized their activity beyond that expected from similar spatial coding, oscillatory firing, and background network activity. When such CA1-mPFC pairs fired together within the same theta cycle, mPFC cells enhanced their phase locking to hippocampal

theta oscillations and maintained a consistent theta phase, even when the phase of their CA1 counterpart varied (Figure 4). These findings suggest a mechanism for interregional assembly interactions that goes beyond coherent oscillatory coupling or firing rate modulation-mediated temporal synchronization.

The role of oscillations in mediating interregional interactions

It has been shown before that theta oscillations synchronize hippocampal and mPFC populations, and such theta oscillation-mediated synchronization can be strengthened at specific task phases.^{25,27,28} While here, theta oscillations provided a reference frame for synchronization, the phase locking of mPFC cells was influenced by the activity of their FnC hippocampal partners. Consequently, the transient increase in the phase locking of mPFC cells signaled an interaction with their CA1 counterparts (Figures 4B–4E). The critical role of the enhanced phase locking in detecting CA1-mPFC functional correlations was demonstrated by the observation that randomizing mPFC theta spike timing during synchronization abolished about 75% of the detected functional correlations and halved the correlation peaks (Figure S7). Furthermore, this mode of interaction did not require the coherent oscillatory firing of the CA1-mPFC cell pair: their

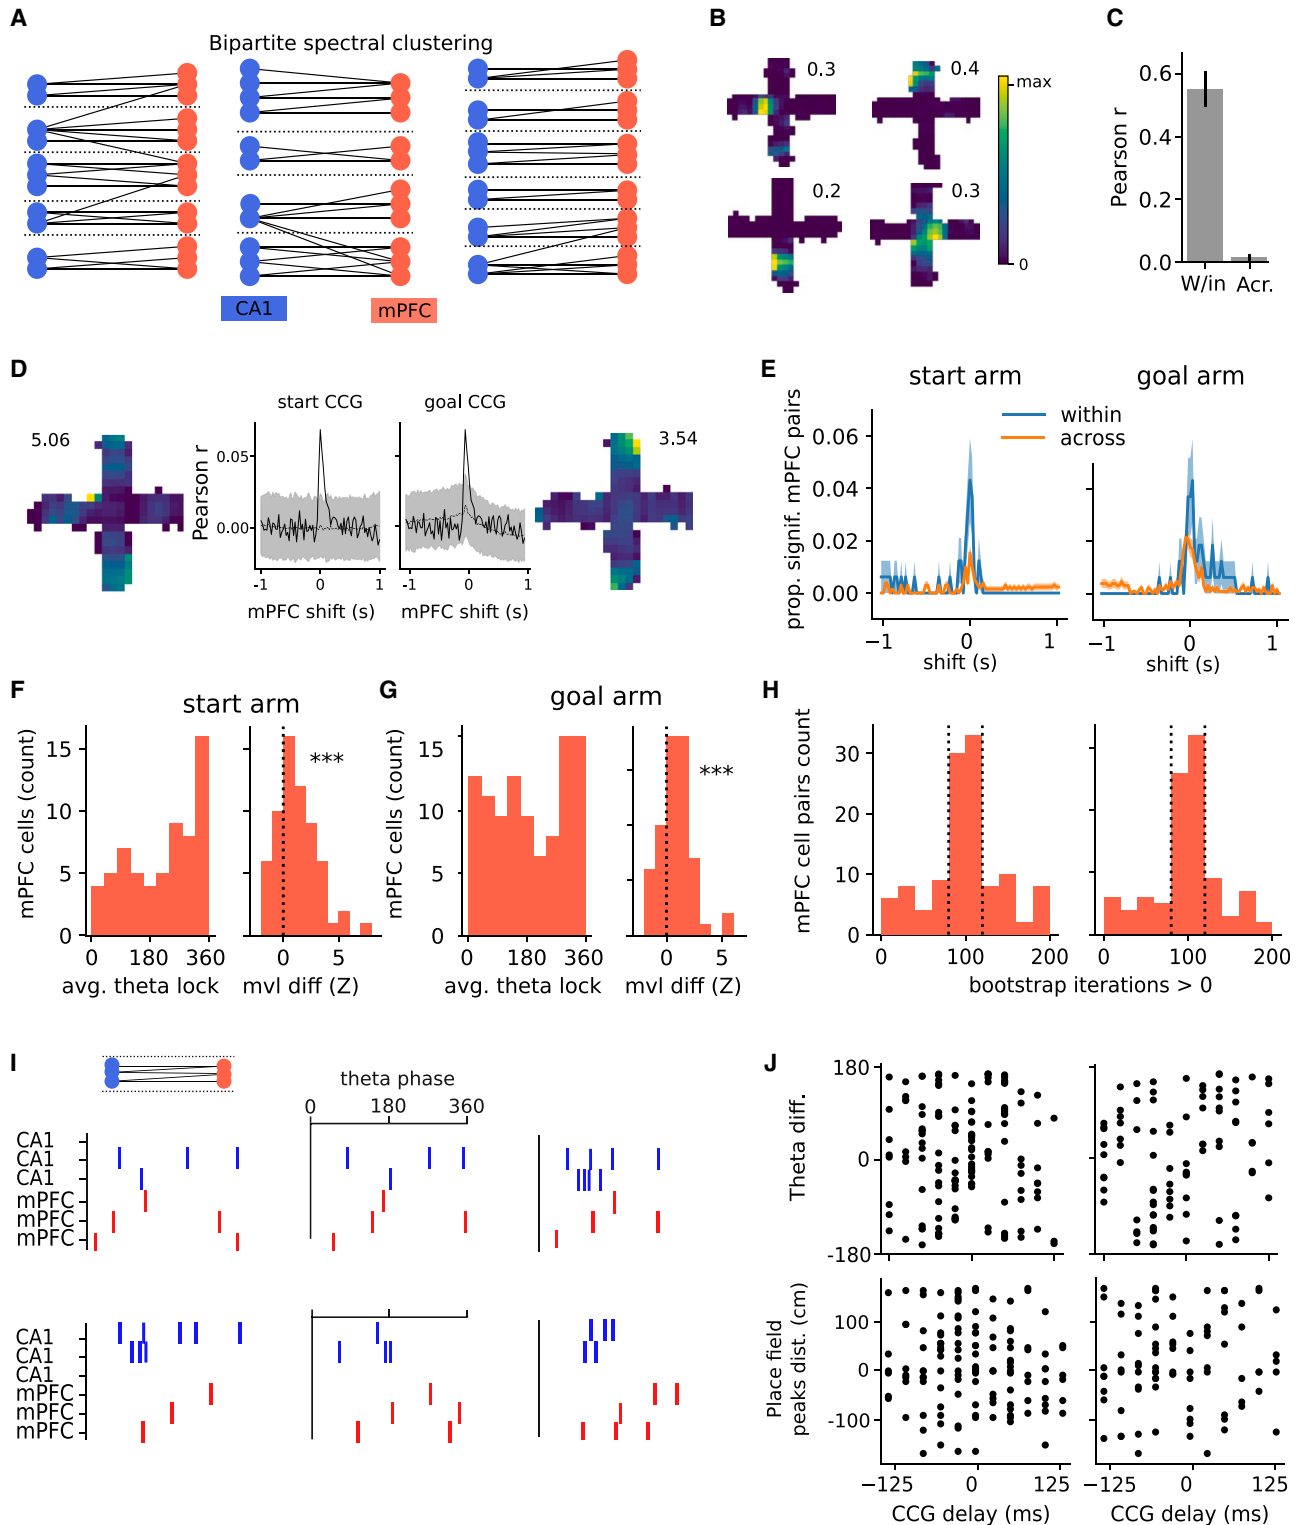


Figure 6. Graph-assembly activity is signaled by mPFC theta sequences

(A) Graph-assembly examples. Functional connections are represented on a graph with vertices representing cells and edges the functional connections. Circles (vertices): CA1 (blue) and mPFC (red) cells. Solid lines (edges): FNCs. Dotted lines: assembly boundaries.

(B) Example graph-assembly firing rate maps.

(C) Graph-assembly rate map similarity between first vs. second half of recordings (W/in) across different assemblies (Acr.). Mann-Whitney U test: $p < 0.001$.

(legend continued on next page)

relative theta spike timing could vary, while the mPFC cell maintained a consistent theta phase (Figures 4F and 4G). In this regard, this mode of synchronization was different from that seen during coherent oscillatory coupling in the visual cortex during sensory binding or attention-related modulation.³⁰ However, gamma-band oscillations can synchronize CA1 and mPFC cells, and such theta-modulated gamma-oscillatory coupling is enhanced between the hippocampus and the mPFC during spatial working memory tasks.³⁶ While we were able to confirm the presence of coherent gamma oscillations during our task as well (Figure 5D), at the times when FnC CA1-mPFC pairs fired together, oscillatory power and coherence in the gamma band did not increase across the CA1 and the mPFC.

In addition to theta-oscillatory coupling, the synchronized firing rate changes of mPFC neurons can also signal interactions across these regions.^{18,20} Such synchronization has been shown during hippocampal SWRs when a subpopulation of mPFC cells increases or decreases their firing.³⁷ The SWR-associated firing of mPFC cells is modulated by behavioral state and context and can signal hippocampal reactivation related to specific arms in a maze.^{17,18,20} Furthermore, in certain conditions, spatially selective mPFC cells can exhibit coherent reactivation with the hippocampus; however, the mPFC can reactivate trajectories independent of the hippocampus.^{15,17} However, in our case, coordinated assemblies did not represent coherent reactivation because we detected FnC CA1-mPFC pairs during active behavior dominated by theta oscillation. Moreover, SWR-associated ripple-band power did not increase in the CA1 region during times when FnC cells fired together (Figure 5). Overall, we showed that across-area interactions were not signaled by changes in the SWR-mediated firing rate nor by coherent gamma-oscillatory firings across the regions. These results all suggest that the transient increase of mPFC phase locking to theta oscillations was the primary factor that marked the functional coupling of specific CA1-mPFC cells and assemblies.

Hippocampal influence on prefrontal spatial coding

Past work suggests that accurate spatial coding requires the mutually interdependent functioning of the intact hippocampus and mPFC circuits.^{19,38–40} Furthermore, one might expect that the execution of our task, which involved rule switching between spatial and non-spatial tasks, requires bidirectional interactions across these regions.⁴¹ Yet, several of our analysis results inde-

pendently suggested a stronger influence of hippocampal assemblies on prefrontal assemblies than in the opposite direction. Firstly, we found that spatial coding in the mPFC lags behind the hippocampus. Coherent location coding and correlated location decoding errors have been shown before between the hippocampus and the mPFC.²³ In our case, both the decoding of locations and decoding errors showed the strongest correspondence when CA1 preceded the mPFC (Figures 1E and 1F). Secondly, independent of spatial coding, a transfer entropy measure also suggested a stronger influence of CA1 population activity on mPFC activity than mPFC influence on CA1 (Figure S1C). Thirdly, as seen before during other spatial tasks,²⁸ mPFC cells exhibited a stronger phase locking to temporally shifted hippocampal theta oscillations that occurred earlier (Figure 3C). Finally, instances of significant functional coupling between CA1-mPFC cells were also more frequent when CA1 cells fired earlier than mPFC cells (Figure 3D).

We also modeled how optimal interactions can be structured to maximize mPFC spatial information when it is derived from CA1 place cell inputs (Figure S4). As we saw for FnC cells in our data, mPFC firing fields in the model were shaped by CA1 place cells that had place fields at an arm where the mPFC cell also possessed a firing field. In agreement with the model, we found in our data that the mPFC firing field can be predicted from the FnC CA1 place fields. Of course, none of these analyses can exclude the possibility that a common input to both regions (such as the MEC) drove parallel hippocampal and mPFC activity, and the delayed mPFC spatial responses were due to more complex mPFC computations that require longer processing times.⁴² Yet, in addition to our analyses, recent work also indicates that CA1 outputs influence mPFC spatial representation: it was shown that the gradient of spatially selective cells in the mPFC correlates with CA1 input density,⁴³ and the intact hippocampus is necessary to maintain precise neocortical spatial representations.⁴⁴

Although, in our case, CA1 responses were consistently earlier during the entire task involving both spatial and non-spatial rules, it is also possible that in certain behavioral conditions, the dominant direction of hippocampus-mPFC communications can be reversed. For example, the correlation (including Granger correlations) of LFPs showed that hippocampal theta oscillations precede mPFC theta oscillations during spatial contextual tasks, while this trend reverses in an object-guided task.⁴⁵ In other cases, the influence may be bidirectional, lasting over several theta cycles.⁴⁶

(D) Example of FnC mPFC pair within the same assembly.

(E) The proportion of significant ($Z > 4.5$) FnC mPFC pairs within vs. across assemblies. Shaded area: 95% CI. $N = 13$ sessions. Mann-Whitney U test: $p < 0.01$ for shifts = -25 , 0 , and $+25$ ms.

(F and G) Histogram of average theta phase (left) and Z scored phase-locking strength (right) of mPFC cells during assembly activity. Phase-locking strength was Z scored relative to bootstrapped non-assembly activity. Binomial tests: all $p < 0.001$.

(H) The number of bootstrap iterations for different assembly cell pairs in which one cell fired at an earlier (i.e., $< 180^\circ$ difference) theta phase than its counterpart during graph-assembly activation. Dotted lines mark the boundary for the pairs in which the preferred theta phases of the pair were significantly different (i.e., one cell fired consistently earlier than the other): 45% start arm and 42% goal arm (binomial test $p < 0.05$).

(I) Example mPFC sequences during graph-assembly activation. Activation patterns of the same graph assembly are displayed, and mPFC cells were ordered according to their theta phase in the top left panel.

(J) Top: the preferred theta-phase difference of mPFC cells during graph-assembly activation vs. cross-correlation (CCG) delay. Left: start arm (Spearman $r = 0.01$, $p = 0.9$). Right: goal arm ($r = 0.19$, $p = 0.058$). Bottom: Place field distance vs. cross-correlation delay. Left: start arm (Spearman $r = -0.09$, $p = 0.302$). Right: goal arm ($r = 0.093$, $p = 0.389$).

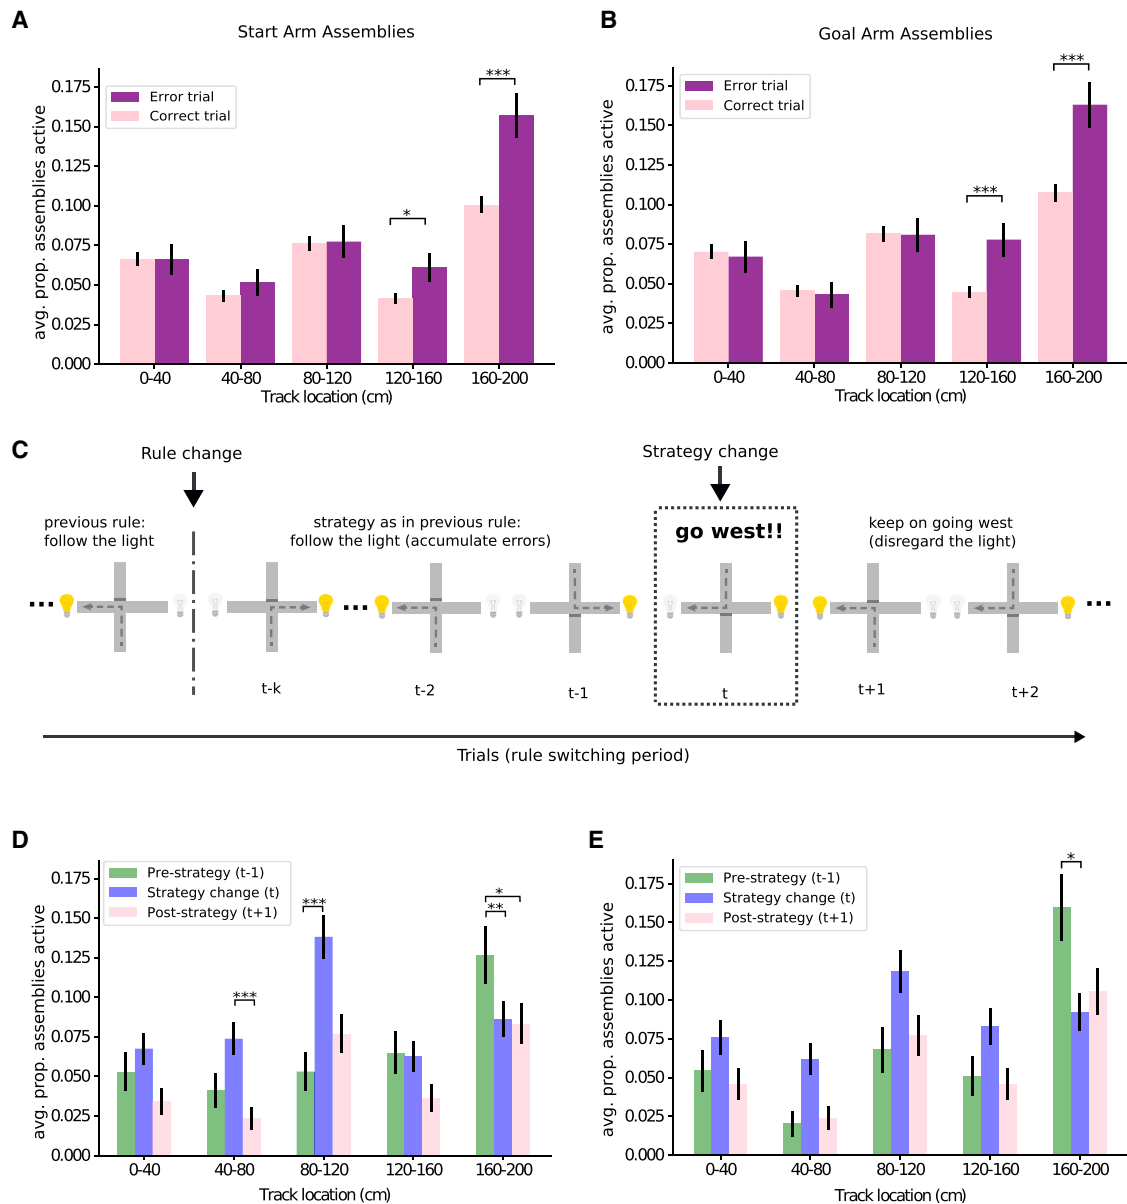


Figure 7. Synchronization of graph assemblies at different behavioral stages

(A and B) The average proportion of active assemblies was compared at correct (N = 719) and error (N = 199) trials at different locations on the maze. Assemblies detected on the start (A) and goal (B) arms were analyzed separately (start assemblies N = 142, goal assemblies N = 133).

(C) Graphical illustration of the rule-switching period used in (D) and (E). This included three trials: the trial when the animal followed the new rule the first time and the trials before and after that.

(D and E) The average proportion of active assemblies was compared before, during, and after rule switching. Assemblies detected at the start (D) and goal (E) locations were analyzed separately (start N = 142, goal N = 133). Error bars: $\pm 95\%$ CI. Three-way ANOVA, Tukey's HSD post hoc, * $p < 0.04$, ** $p < 0.01$, *** $p = 0.001$.

Prefrontal delay as a function of cognitive load

We also found that the delay between the CA1 and the mPFC was shorter when cognitive demand was higher in the start arm while the animal had to choose the goal arm. We found this in relation to mPFC phase locking to hippocampal theta oscillations and for the significant CA1-mPFC functional correlations (Figures 3C and 3D). This may indicate that CA1-mPFC in-

teractions were enhanced when cognitive demand was high before a spatial decision had to be performed. This is in line with studies on theta-oscillatory synchronization that showed enhanced coupling before the choice point or when rules have changed.^{25,27} It is possible that the routing of spatial information follows different paths at different processing stages or that the mPFC cells are tuned to respond to hippocampal inputs with a

shorter delay at the start arm. Note that in previous work, hippocampal-driven mPFC activity was seen only during the encoding but not the recall stage of a spatial working memory task,⁴⁷ indicating that the effectiveness of the hippocampal spatial input to the mPFC can be modulated. Nevertheless, in our case, the hippocampal drive was maintained during the entire task but with different delays.

It is also possible that the routing of spatial information follows different paths with different types of cognitive load. In support of this observation is the fact that direct projections from the ventral hippocampus to the mPFC are important only for task periods requiring memory encoding.⁴⁷ There are several possible routes via which the hippocampus could transmit spatial information to the mPFC. Apart from the prominent direct projections arising from the ventral hippocampus, indirect projections can be established through entorhinal and perirhinal cortices.^{48,49} Information may also be transferred indirectly via the thalamic nucleus reuniens; this area bidirectionally connects with the mPFC,⁵⁰ projects to the hippocampus,⁵⁰ and in turn receives hippocampal input mostly via the subiculum.⁵¹

Interacting assemblies can use different temporal coding schemes

Here, we also showed that FnC CA1-mPFC pairs could be subdivided into separate assemblies, and the theta spike timing of mPFC cells can signal the activation of specific CA1-mPFC interregional assemblies (Figure 6). Because different mPFC cells of an assembly showed different theta-phase preferences during assembly activation, we expect that mPFC cells fire in theta sequences during assembly activation. However, the order in which mPFC cells fired during assembly activation did not reflect their temporal activation pattern on the maze. Therefore, these sequences did not represent reactivated behavioral trajectories or theta sequences, which occur in other instances.^{15,23} As a consequence, different temporal coding schemes can operate in the mPFC: while spatially selective mPFC cells can transiently reactivate trajectories, at other times, different sequential activation patterns can indicate the activity of specific CA1 assemblies. Furthermore, sequential activity patterns that represent trajectories are different from those related to interregional assemblies; hence, the two coding schemes do not interfere with each other and can be differentiated and decoded independently in downstream areas.

Our findings also suggest that different temporal coding schemes can operate in the hippocampus and in the mPFC during the interactions, considering that the relative spike timing of CA1-mPFC pairs could vary within a theta cycle. While the theta spike timing of CA1 cells can code for place because of theta-phase precession,²⁹ or they can express out-of-place theta sequences,^{52,53} mPFC theta phases did not follow the same pattern during periods of interactions. Similar interaction schemes may apply to other brain areas as well, considering that theta oscillations are seen in widespread brain regions beyond the hippocampal formation.⁵⁴ Theta-related assembly synchronization can enable multiple regions to function and interact with each other in the limbic system and beyond. Our data also indicate that assembly interactions through theta syn-

chronization can allow each region to use its own unique temporal coding scheme during the time of interactions. Temporal coding may not be simply transferred but often transformed across different cortical processing stages, necessitating further experimental and theoretical work to understand how temporal codes in one region are decoded and transformed downstream. Overall, our work suggests that theta oscillations play crucial roles in organizing interregional assemblies.

Limitations of the study

The work relied on a GLM model to account for possible common factors that may cause the correlated firing of the CA1-mPFC cell pairs. Of course, we could not account for possible common factors that are currently unknown. That said, we were able to show that the excess correlation (i.e., functional coupling that cannot be explained by the GLM) we observed was largely due to the increased theta phase locking of the mPFC cells when they fired together with their CA1 counterparts. We could not specify a physiological mechanism that underlies such synchronization-dependent oscillatory phase-locking enhancement, and we could only speculate about the anatomical pathways that are involved in it. Finally, we have limited data investigating the possible behavioral role of the activity of CA1-mPFC assemblies. Nonetheless, we show that FnC CA1-mPFC graph assemblies were activated more strongly at error trials; thus, in the current behavioral context, they may signal the need for changing the rule of the task.

STAR★METHODS

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

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
 - Lead contact
 - Materials availability
 - Data and code availability
- EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS
- METHOD DETAILS
 - Subjects and surgery
 - Plus maze apparatus and task
 - Beginning of good performance
 - Histology and reconstruction of recording positions
- QUANTIFICATION AND STATISTICAL ANALYSIS
 - Data acquisition
 - Spike sorting and unit classification
 - Dataset details
 - Statistical analysis
 - Modeling
 - Programming languages and packages used

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.celrep.2023.113015>.

ACKNOWLEDGMENTS

We thank A. Cumpelik, H. Chioffi, and L. Bollman for comments on an earlier version of this manuscript. This work was funded by EU-FP7 MC-ITN IN-SENS (grant 607616).

AUTHOR CONTRIBUTIONS

M.N., K.K., and J.C. conceptualized the study. K.K. conducted the experiments, M.N., K.K., and F.S. analyzed the data. M.N. and J.C. wrote the text, and all authors contributed to manuscript preparation.

DECLARATION OF INTERESTS

The authors declare no competing interests.

INCLUSION AND DIVERSITY

We support inclusive, diverse, and equitable conduct of research.

Received: September 14, 2022

Revised: June 21, 2023

Accepted: August 4, 2023

REFERENCES

- Hebb, D.O. (1949). *The Organization of Behavior* (Wiley).
- Buzsáki, G. (2010). Neural syntax: cell assemblies, synapse ensembles, and readers. *Neuron* 68, 362–385. <https://doi.org/10.1016/j.neuron.2010.09.023>.
- Buzsáki, G., Bickford, R.G., Ponomareff, G., Thal, L.J., Mandel, R., and Gage, F.H. (1988). Nucleus basalis and thalamic control of neocortical activity in the freely moving rat. *J. Neurosci.* 8, 4007–4026. <https://doi.org/10.1523/JNEUROSCI.08-11-04007.1988>.
- Averbeck, B.B., Latham, P.E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* 7, 358–366. <https://doi.org/10.1038/nrn1888>.
- Cohen, M.R., and Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nat. Neurosci.* 14, 811–819. <https://doi.org/10.1038/nn.2842>.
- Newsome, W.T., Britten, K.H., and Movshon, J.A. (1989). Neuronal correlates of a perceptual decision. *Nature* 341, 52–54. <https://doi.org/10.1038/341052a0>.
- Harris, K.D., Csicsvari, J., Hirase, H., Dragoi, G., and Buzsáki, G. (2003). Organization of cell assemblies in the hippocampus. *Nature* 424, 552–556. <https://doi.org/10.1038/nature01834>.
- Abeles, M. (1982). *Local Cortical Circuits* (Springer).
- Ylinen, A., Bragin, A., Nádasdy, Z., Jandó, G., Szabó, I., Sik, A., and Buzsáki, G. (1995). Sharp wave-associated high-frequency oscillation (200 Hz) in the intact hippocampus: network and intracellular mechanisms. *J. Neurosci.* 15, 30–46. <https://doi.org/10.1523/JNEUROSCI.15-01-00030.1995>.
- Khodagholy, D., Gelineas, J.N., and Buzsáki, G. (2017). Learning-enhanced coupling between ripple oscillations in association cortices and hippocampus. *Science* 358, 369–372. <https://doi.org/10.1126/science.aan6203>.
- Logothetis, N.K., Eschenko, O., Murayama, Y., Augath, M., Stuedel, T., Evrard, H.C., Besserve, M., and Oeltermann, A. (2012). Hippocampal-cortical interaction during periods of subcortical silence. *Nature* 491, 547–553. <https://doi.org/10.1038/nature11618>.
- Ji, D., and Wilson, M.A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat. Neurosci.* 10, 100–107. <https://doi.org/10.1038/nn1825>.
- Olafsdottir, H.F., Carpenter, F., and Barry, C. (2016). Coordinated grid and place cell replay during rest. *Nat. Neurosci.* 19, 792–794. <https://doi.org/10.1038/nn.4291>.
- 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. <https://doi.org/10.1016/j.neuron.2019.09.012>.
- Kaefer, K., Nardin, M., Blahna, K., and Csicsvari, J. (2020). Replay of Behavioral Sequences in the Medial Prefrontal Cortex during Rule Switching. *Neuron* 106, 154–165.e6. <https://doi.org/10.1016/j.neuron.2020.01.015>.
- O'Neill, J., Boccarda, C.N., Stella, F., Schoenenberger, P., and Csicsvari, J. (2017). Superficial layers of the medial entorhinal cortex replay independently of the hippocampus. *Science* (New York, N.Y.) 355, 184–188. <https://doi.org/10.1126/science.aag2787>.
- Tang, W., Shin, J.D., and Jadhav, S.P. (2021). Multiple time-scales of decision-making in the hippocampus and prefrontal cortex. *Elife* 10, e66227. <https://doi.org/10.7554/eLife.66227>.
- Berners-Lee, A., Wu, X., and Foster, D.J. (2021). Prefrontal Cortical Neurons Are Selective for Non-Local Hippocampal Representations during Replay and Behavior. *J. Neurosci.* 41, 5894–5908. <https://doi.org/10.1523/JNEUROSCI.1158-20.2021>.
- Patai, E.Z., and Spiers, H.J. (2021). The Versatile Wayfinder: Prefrontal Contributions to Spatial Navigation. *Trends Cognit. Sci.* 25, 520–533. <https://doi.org/10.1016/j.tics.2021.02.010>.
- 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. <https://doi.org/10.1038/s41467-018-04498-w>.
- 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. <https://doi.org/10.1038/nn.2134>.
- Guise, K.G., and Shapiro, M.L. (2017). Medial Prefrontal Cortex Reduces Memory Interference by Modifying Hippocampal Encoding. *Neuron* 94, 183–192.e8. <https://doi.org/10.1016/j.neuron.2017.03.011>.
- Zielinski, M.C., Shin, J.D., and Jadhav, S.P. (2019). Coherent coding of spatial information mediated by theta oscillations in the hippocampus and prefrontal cortex. *J. Neurosci.* 39, 4550–4565. <https://doi.org/10.1523/jneurosci.0106-19.2019>.
- Bossomaier, T., Barnett, L., Harré, M., and Lizier, J.T. (2016). *An Introduction to Transfer Entropy* (Cham: Springer International Publishing), pp. 65–95.
- Jones, M.W., and Wilson, M.A. (2005). Theta Rhythms Coordinate Hippocampal-Prefrontal Interactions in a Spatial Memory Task. *PLoS Biol.* 3, e402. <https://doi.org/10.1371/journal.pbio.0030402>.
- Hardcastle, K., Ganguli, S., and Giocomo, L.M. (2015). Environmental Boundaries as an Error Correction Mechanism for Grid Cells. *Neuron* 86, 827–839. <https://doi.org/10.1016/j.neuron.2015.03.039>.
- 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. <https://doi.org/10.1016/j.neuron.2010.05.013>.
- Siapas, A.G., Lubenov, E.V., and Wilson, M.A. (2005). Prefrontal Phase Locking to Hippocampal Theta Oscillations. *Neuron* 46, 141–151. <https://doi.org/10.1016/j.neuron.2005.02.028>.
- O'Keefe, J., and Recce, M.L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3, 317–330. <https://doi.org/10.1002/hipo.450030307>.
- Fries, P., Nikolić, D., and Singer, W. (2007). The gamma cycle. *Trends Neurosci.* 30, 309–316. <https://doi.org/10.1016/j.tins.2007.05.005>.

31. Colgin, L.L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M.B., and Moser, E.I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* 462, 353–357. <https://doi.org/10.1038/nature08573>.
32. Newman, M.E.J. (2006). Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E* 74, 036104. <https://doi.org/10.1103/PhysRevE.74.036104>.
33. Pothén, A., Simon, H.D., and Liou, K.-P. (1990). Partitioning Sparse Matrices with Eigenvectors of Graphs. *SIAM J. Matrix Anal. Appl.* 11, 430–452. <https://doi.org/10.1137/0611030>.
34. Dragoi, G., and Buzsáki, G. (2006). Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron* 50, 145–157. <https://doi.org/10.1016/j.neuron.2006.02.023>.
35. Skaggs, W.E., McNaughton, B.L., Wilson, M.A., and Barnes, C.A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6, 149–172. [https://doi.org/10.1002/\(SICI\)1098-1063\(1996\)6:2<149::AID-HIPO6>3.0.CO;2-K](https://doi.org/10.1002/(SICI)1098-1063(1996)6:2<149::AID-HIPO6>3.0.CO;2-K).
36. Tamura, M., Spellman, T.J., Rosen, A.M., Gogos, J.A., and Gordon, J.A. (2017). Hippocampal-prefrontal theta-gamma coupling during performance of a spatial working memory task. *Nat. Commun.* 8, 2182. <https://doi.org/10.1038/s41467-017-02108-9>.
37. 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. <https://doi.org/10.1016/j.neuron.2016.02.010>.
38. Hok, V., Chah, E., Save, E., and Poucet, B. (2013). Prefrontal Cortex Focally Modulates Hippocampal Place Cell Firing Patterns. *J. Neurosci.* 33, 3443–3451. <https://doi.org/10.1523/jneurosci.3427-12.2013>.
39. Ito, H.T., Zhang, S.-J., Witter, M.P., Moser, E.I., and Moser, M.-B. (2015). A prefrontal-thalamo-hippocampal circuit for goal-directed spatial navigation. *Nature* 522, 50–55. <https://doi.org/10.1038/nature14396>.
40. Navawongse, R., and Eichenbaum, H. (2013). Distinct Pathways for Rule-Based Retrieval and Spatial Mapping of Memory Representations in Hippocampal Neurons. *J. Neurosci.* 33, 1002–1013. <https://doi.org/10.1523/jneurosci.3891-12.2013>.
41. Malagon-Vina, H., Ciocchi, S., Passecker, J., Dorffner, G., and Klausberger, T. (2018). Fluid network dynamics in the prefrontal cortex during multiple strategy switching. *Nat. Commun.* 9, 309. <https://doi.org/10.1038/s41467-017-02764-x>.
42. Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84. <https://doi.org/10.1038/nature12742>.
43. Sauer, J.-F., Folschweiller, S., and Bartos, M. (2022). Topographically organized representation of space and context in the medial prefrontal cortex. *Proc. Natl. Acad. Sci. USA* 119, e2117300119. <https://doi.org/10.1073/pnas.2117300119>.
44. Esteves, I.M., Chang, H., Neumann, A.R., Sun, J., Mohajerani, M.H., and McNaughton, B.L. (2021). Spatial Information Encoding Across Multiple Neocortical Regions Depends on an Intact Hippocampus. *J. Neurosci.* 41, 307–319. <https://doi.org/10.1523/jneurosci.1788-20.2020>.
45. Place, R., Farovik, A., Brockmann, M., and Eichenbaum, H. (2016). Bidirectional prefrontal-hippocampal interactions support context-guided memory. *Nat. Neurosci.* 19, 992–994. <https://doi.org/10.1038/nn.4327>.
46. Yu, J.Y., and Frank, L.M. (2021). Prefrontal cortical activity predicts the occurrence of nonlocal hippocampal representations during spatial navigation. *PLoS Biol.* 19, e3001393. <https://doi.org/10.1371/journal.pbio.3001393>.
47. Spellman, T., Rigotti, M., Ahmari, S.E., Fusi, S., Gogos, J.A., and Gordon, J.A. (2015). Hippocampal-prefrontal input supports spatial encoding in working memory. *Nature* 522, 309–314. <https://doi.org/10.1038/nature14445>.
48. Delatour, B., and Witter, M.P. (2002). Projections from the parahippocampal region to the prefrontal cortex in the rat: evidence of multiple pathways. *Eur. J. Neurosci.* 15, 1400–1407. <https://doi.org/10.1046/j.1460-9568.2002.01973.x>.
49. Hoover, W.B., and Vertes, R.P. (2007). Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. *Brain Struct. Funct.* 212, 149–179. <https://doi.org/10.1007/s00429-007-0150-4>.
50. Vertes, R.P. (2006). Interactions among the medial prefrontal cortex, hippocampus and midline thalamus in emotional and cognitive processing in the rat. *Neuroscience* 142, 1–20. <https://doi.org/10.1016/j.neuroscience.2006.06.027>.
51. Varela, C., Kumar, S., Yang, J.Y., and Wilson, M.A. (2014). Anatomical substrates for direct interactions between hippocampus, medial prefrontal cortex, and the thalamic nucleus reuniens. *Brain Struct. Funct.* 219, 911–929. <https://doi.org/10.1007/s00429-013-0543-5>.
52. Feng, T., Silva, D., and Foster, D.J. (2015). Dissociation between the Experience-Dependent Development of Hippocampal Theta Sequences and Single-Trial Phase Precession. *J. Neurosci.* 35, 4890–4902. <https://doi.org/10.1523/JNEUROSCI.2614-14.2015>.
53. 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. <https://doi.org/10.1016/j.neuron.2010.01.034>.
54. Buzsáki, G. (2002). Theta oscillations in the hippocampus. *Neuron* 33, 325–340. [https://doi.org/10.1016/s0896-6273\(02\)00586-x](https://doi.org/10.1016/s0896-6273(02)00586-x).
55. 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. [https://doi.org/10.1016/s0896-6273\(00\)80525-5](https://doi.org/10.1016/s0896-6273(00)80525-5).
56. 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. <https://doi.org/10.1152/jn.2000.84.1.401>.
57. Csicsvari, J., Hirase, H., Czurko, A., Mamiya, A., and Buzsáki, G. (1999). Oscillatory Coupling of Hippocampal Pyramidal Cells and Interneurons in the Behaving Rat. *J. Neurosci.* 19, 274–287. <https://doi.org/10.1523/jneurosci.19-01-00274.1999>.
58. 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. <https://doi.org/10.1152/jn.1998.79.2.1017>.
59. Priesemann, V., Wibral, M., Valderrama, M., Pröpper, R., Le Van Quyen, M., Geisel, T., Triesch, J., Nikolić, D., and Munk, M.H.J. (2014). Spike avalanches in vivo suggest a driven, slightly subcritical brain state. *Front. Syst. Neurosci.* 8, 108. <https://doi.org/10.3389/fnsys.2014.00108>.
60. Lizier, J.T. (2014). JIDT: An Information-Theoretic Toolkit for Studying the Dynamics of Complex Systems. *Front. Robot. AI* 1.
61. Boccara, C.N., Nardin, M., Stella, F., O'Neill, J., and Csicsvari, J. (2019). The entorhinal cognitive map is attracted to goals. *Science* 363, 1443–1447. <https://doi.org/10.1126/science.aav4837>.
62. Agresti, A. (2015). *Foundations of Linear and Generalized Linear Models* (John Wiley & Sons Inc).
63. Nardin, M., Csicsvari, J., Tkačik, G., and Savin, C. (2021). The Structure of Hippocampal CA1 Interactions Optimizes Spatial Coding across Experience. <https://doi.org/10.1101/2021.09.28.460602>.
64. Skaggs, W.E., McNaughton, B.L., and Gothard, K.M. (1993). An information-theoretic approach to deciphering the hippocampal code. In *Advances in Neural Information Processing Systems*, pp. 1030–1037. <https://doi.org/10.1523/ENEURO.0266-21.2021>.
65. Souza, B.C., Pavão, R., Belchior, H., and Tort, A.B.L. (2018). On Information Metrics for Spatial Coding. *Neuroscience* 375, 62–73. <https://doi.org/10.1016/j.neuroscience.2018.01.066>.

66. Welch, P. (1967). The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Trans. Audio Electroacoust.* 15, 70–73. <https://doi.org/10.1109/TAU.1967.1161901>.
67. Barber, M.J. (2007). Modularity and community detection in bipartite networks. *Phys. Rev. E - Stat. Nonlinear Soft Matter Phys.* 76, 066102. <https://doi.org/10.1103/PhysRevE.76.066102>.
68. McCullagh, P., and Nelder, J.A. (1989). *Generalized Linear Models* (Springer US). <https://doi.org/10.1007/978-1-4899-3242-6>.
69. Smolensky, P. (1986). Information processing in dynamical systems: foundations of harmony theory. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (MIT Press), pp. 194–281. foundations.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental models: Organisms/strains		
Long Evans Rats	Janvier, France	RjOr:LE
Software and algorithms		
KlustaKwik	(Harris et al., 2000) ⁵⁶	https://github.com/klusta-team/klustakwik/
Python	Python Software Foundation	https://www.python.org
R	R Foundation	https://www.r-project.org/
Original code	This paper	https://doi.org/10.5281/zenodo.8193614
Other		
12 μ m tungsten wires	California Fine Wire	M294520
Headstage amplifier	Axona, St. Albans, UK	http://www.axona.com/

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources and reagents should be directed to and fulfilled by the lead contact, Jozsef Csicsvari (jozsef.csicsvari@ist.ac.at).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- Data used in this study will be made available upon request by contacting the lead contact, Jozsef Csicsvari (jozsef.csicsvari@ist.ac.at).
- All original code is available in a GitHub repository: <https://doi.org/10.5281/zenodo.8193614>
- Any additional information required to reanalyze the data reported in this work paper is available from the corresponding authors upon request.

EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

Four male Long-Evans rats (300–350 g, 2–4 months of age; Janvier, France) were used in the study. The animals were housed in a separate room on a 12 h 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: BMFW-66.018/0015-WF/V3b/2014).

METHOD DETAILS

The data used in this study is the same as used in.¹⁵ We will report the experimental methods here for completeness.

Subjects and surgery

Rats were implanted with microdrives housing 32 individually-movable tetrodes, arranged into three bundles targeting the right dorsal hippocampus (specifically dorsal CA1, hippocampus) and left and right medial prefrontal cortex (specifically prelimbic area, mPFC). The CA1 bundle consisted of 16 tetrodes and the two mPFC bundles of 8 tetrodes each. Tetrodes were fabricated out of four 12 mm 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 CA1 bundle. The tips of the tetrodes were gold-plated to reduce the impedance to around 300 k Ω . 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 hippocampus (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 were centered above their respective craniotomies and lowered into the brain at a depth of 2 mm for the mPFC and 1 mm for the hippocampus. 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 of recovery. Thereafter, tetrodes were gradually moved in 50–200 μm steps into the CA1 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, Richmond, 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 the 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 5s passed, the animal was manually picked up and placed in the rest box before commencing the next trial after a delay of 10s. 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 the reward. To prevent an odor-guided strategy, pellet dust was scattered along the maze and pellet-filled cups were invisible to the animal that was placed under both goal arms. On each recording day, the animal underwent behavioral blocks as follows: rest, rule 1 (previous day's old rule), rest, pre-switch, rule switching, post-switch, rest, rule 2 (new rule), rest.

After the first rest, the animal started by performing trials based on the previous day's old rule. After reaching the performance criterion (see below), the animal rested again and afterward, the rule-switching phase began.

During the pre-switch block, the animal had to collect a reward based on the last rule of the previous day until reaching the performance criterion (see below). Then the rule was changed and the reward had to be collected based on the new rule. The change in the 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 rule-switching block, while the post-switch block comprised all trials from the beginning of good performance (defined in the next sub-subsection). 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 the correct performance of a spatial rule involves two trajectories (e.g., go-east rule: north to east and south to east), the 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 a similar number of light rule trials where the animal performed trajectories that matched those of the spatial rule.

After another rest session, the animal performed a final 20 trials of the newly acquired rule.

Beginning of good performance

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.

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

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 (300 mg/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 subsections with a cryostat (50–60 μm), mounted on glass slides and stained with cresyl violet. The positions of tetrode tips were determined from stained subsections and cells recorded from tetrodes outside mPFC were excluded from analysis. For cells recorded from CA1 tetrodes the presence of SWRs in the field recordings served as inclusion criteria.

QUANTIFICATION AND STATISTICAL ANALYSIS

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 x,y location of the animal. Every day before recording, CA1 tetrodes were moved to optimize the yield of recorded cells. Additionally, mPFC tetrodes were lowered every day by 30–50 mm to ensure the recording of a new population of cells.

Spike sorting and unit classification

Clustering of spikes and unit isolation procedures were described previously.⁵⁵ The raw data were resampled 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>).⁵⁶ 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.⁵⁶ Putative excitatory principal cells and inhibitory interneurons were discriminated using their auto-correlograms, firing rates and waveforms.⁵⁷

Dataset details

Across the four animals, we recorded, processed, and analyzed data from a total of 13 experimental recording days (4,4,3,2 respectively).

In all our analysis we included only cells with an average firing rate >0.25 spikes per second in each of the three experimental phases. This comprised a total of 530 hippocampal (putative) principal cells and 160 interneurons, and 477 prefrontal (putative) principal cells and 105 interneurons.

Statistical analysis

Linearized position

To linearize the behavior of the animal, we calculated the distance to the center from the 2D spatial position of the animal.¹⁵ 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.

Firing rate maps

We inferred the average firing of each cell at each given location of the environment separately for three cases: 2D maps, 1D trajectory dependent, and 1D trajectory independent. In every case, we utilized data from periods when the animal was moving faster than 7 cm/s to avoid potential nonlocal population activity.¹⁵

- 2D maps: We binned the x,y locations of the animal in 5cm square bins. As explained in the experimental setup, spatial coordinates were acquired every $dt = 25.6ms$.

We counted how much time (in seconds) the animal spent in each location, which corresponded to the occupancy map.

Denoting with $\hat{x}(t), \hat{y}(t)$ the coordinates of the animal at time t , we have

$$occ(x,y) = \sum_{t=1}^T \frac{(\hat{x}(t) = x) \cdot (\hat{y}(t) = y)}{dt}$$

where $a = b$ is 1 if and only if $a = b$, and is 0 otherwise.

Afterward, for each cell i , we counted the number of spikes emitted in each location, and divided by the time spent there. In formulas, denoting by t_1^i, \dots, t_K^i the times at which K spikes were emitted, we define the rate

$$\hat{\lambda}(x,y) = \sum_{t=t_1^i}^{t_K^i} \frac{(\hat{x}(t) = x) \cdot (\hat{y}(t) = y)}{occ(x,y)}$$

Finally, we regularized $\hat{\lambda}(x,y)$ by convolving with a symmetric 2D Gaussian kernel $\kappa(h,k)$ with $\sigma = 2$ bins and truncated at a maximum width $w = 10$ bins:

$$\lambda(x,y) = \sum_{h=-w}^w \sum_{k=-w}^w \hat{\lambda}(x+h,y+k) \kappa(h,k)$$

Locations x,y with zero occupancy were excluded from the convolution. Locations near the border used only the portion of the kernel that would not exceed the edges (portion appropriately rescaled so as to sum up to 1).

- Trajectory-dependent 1D maps: Since the arms of the maze are relatively narrow, we computed a linearized version of the firing rate maps, effectively yielding an average firing as a function of the distance to the goal, separately for the 4 trajectories: North to East (NE), North to West (NW), South to East (SE), South to West (SW) (Figure 1A). In words, we selected the trials where the animal followed one trajectory, and with that data binned the linearized position in 10 cm bins and counted the time spent at each discrete location. Afterward, for each cell and each trajectory, we counted the number of spikes emitted in each binned location and divided it by the time spent there. Finally, we convolved the rate maps of each cell for each trajectory with a 1D Gaussian kernel with $\sigma = 1$ bin.
- Trajectory independent 1D maps: We computed the same quantity as above, without separating the trials into 4 trajectory groups.

Population vector similarity

Population vector (PV) similarity is a measure that allows us to quantify the change in average population activity at any given location across contexts.

The PV at a particular location x represents the vector of the average activities of all the cells in the population under study:

$$PV(x) = \{\lambda_1(x), \dots, \lambda_N(x)\}$$

We constructed them starting from the previously computed firing rate maps, separately for the two populations and the two rules. We employed a Pearson correlation to quantify the similarity of the PV for each location across rules.

Symmetry of firing rate maps

We measured the “mirroredness” of firing rates to check whether single cells fire similarly at opposite sides of the track. To do that, we measured the Pearson correlation of the linearized firing rate maps of opposite arms (S vs. N and E vs. W) and computed the average of the two measures.

Afterward, we computed a null-distribution of symmetry measures obtained from 200 random wrapping of the linearized rate maps (for each cell individually). We deemed a rate map symmetric if its symmetry score was higher than the 95th percentile of the corresponding null-distribution.

The random wrapping of rate maps worked as follows. Consider a rate map $\lambda = \{\lambda(0), \dots, \lambda(M)\}$, ordered as usual from location 0 to location M . We randomly selected an integer $k \in \{1, \dots, M - 1\}$, and considered the shuffled rate map $\lambda_{shuf} = \{\lambda(k), \dots, \lambda(M), \lambda(0), \dots, \lambda(M - k)\}$

Decoding of distance to the goal

For each decoding scheme, we randomly selected 75% of the trials for computing the firing rate maps, and used the remaining 25% for decoding and assessing decoding quality. We employed a maximum a posteriori (MAP) Bayesian decoding under the assumption that cells are independent and fire according to a Poisson distribution.⁵⁸

Denoting the firing rate (measured in spikes per second) of cell i at location x with $\lambda_i(x)$, and the same cell emitted n_i spikes in a period dt of time (measured in seconds), the probability that the animal is at location x given the spiking activity is:

$$P(x | n_i) \propto P(x) P(n_i | x) = P(x) e^{-dt \lambda_i(x)} \frac{(dt \lambda_i(x))^{n_i}}{n_i!}$$

Combining the information from multiple cells $i = 1, \dots, N$, and assuming independence:

$$P(x | \{n_1, \dots, n_N\}) \propto \prod_{i=1}^N P(x | n_i)$$

Assuming a uniform prior $P(x)$, we defined our decoded location x_{MAP} as

$$x_{MAP} = \underset{x}{\operatorname{argmax}} P(x | \{n_1, \dots, n_N\})$$

Correlation of decoding errors

To check whether the errors of one area agree with the errors of the other area and whether this effect increases at a delay, we employed a delayed version of an analysis proposed by Zielinski and colleagues.²³

If the spatial encoding of PFC is delayed compared to CA1, in that either PFC receives information from CA1, or that information takes longer to reach PFC, then spikes in PFC should represent spatial information that is older than CA1's. Hence, we took all mPFC spiking times and added a time lag $\tau \in [-500, 500]$ ms, and for each lag we computed 1D firing rate maps on a fraction of data (random 75%) and decoded the position from the remaining part of the data. Afterward, for each delay, we computed the Spearman correlation between 1D decoding errors from CA1 (without lag) and the lagged PFC activity.

Transfer entropy

Transfer entropy is a non-parametric measure of directed (time-asymmetric) transfer of information between two random processes.⁵⁹ Transfer entropy from a process X to another process Y is the amount of uncertainty reduced in future values of Y by knowing the past values of X , conditioned on past values of Y .

More specifically, if X_t, Y_t denote two random processes, transfer entropy from X_t to Y_t is defined as the conditional mutual information between Y_t and the history of X_t , denoted by $X_{t-1, t-2, \dots}$, conditioned on the history of the influenced variable $Y_{t-1, t-2, \dots}$:

$$T_{X \rightarrow Y} = MI(Y_t, X_{t-1, t-2, \dots} | Y_{t-1, t-2, \dots})$$

We used the JIDT package⁶⁰ to estimate this quantity, which was measured across binned (25.6ms) and binarized spike trains of each pair of CA1-PFC cells. Significance was measured by comparing the actual value against 1000 values obtained by randomly shifting the spike trains (uniform random from 1 to 100 s).⁶¹

Detection of cross-area couplings

We employed a statistical modeling approach to detect pairs of CA1-PFC cells that are significantly functionally correlated. With “functionally correlated” we denote cell pairs across brain areas whose noise correlation⁵ is much higher than expected from our null-models. For each cell, we fitted a generalized linear model (GLM⁶² that included all possible covariates measured which could influence and explain the cross-area correlations. These covariates were: linearized spatial position, trajectory, theta selectivity, speed selectivity, spiking history, and within-area spiking of other cells (i.e., PFC cells were fitted with the spiking of other PFC cells only and, separately, CA1 cells with the spiking of the other CA1 cells only; for further details see subsection “Modelling” below).

These models were used to compute a statistical test, the null hypothesis being that the cross-correlogram among cell pairs is completely explained by external covariates. The alternative hypothesis is that external covariates cannot explain the amount of co-variability; in this case, we considered the correlation as being functional. This approach is similar to the one introduced in.⁶³

With these GLM null-models, we simulated the activity of each cell 10000 times and, for each CA1-PFC cell pair, a cross-correlogram of the responses was computed. Those surrogate cross-correlograms were used to measure how much the actual pairwise cross-correlation measured on real data differed from the simulated ones. We did this for each possible PFC delay in the range of + -1 s.

We considered a cell pair significantly coupled if the peak within -125, ..., +125 ms of the actual cross-correlogram exceeded the mean plus 4.5 standard deviations of the peaks of the 10000 surrogate cross-correlograms for that pair. This threshold is very conservative and was chosen to compensate for multiple comparisons, representing a Bonferroni correction.

Cells in one area that showed significant functional correlation to at least one cell in the other area were termed functionally coupled, often abbreviated to F.C. or FnC in our figures.

Complementarity measure

To measure if two cells are complementary, i.e., if their joint activity could help to form a symmetric pattern, we computed the Pearson correlation of the linearized firing rate maps for opposite trajectories (for example, SW for cell 1 vs. NE for cell 2) and then reported the average across the possible comparisons (SW vs. NE and SE vs. NW).

mPFC rate map prediction

We predicted the firing rate maps of mPFC cells starting from the firing rate maps of (simultaneously recorded) CA1 cells. To do so we utilized only functionally coupled cell pairs (see definition above). We summed CA1 maps and weighted them by using the strength of their peak functional correlation (i.e., peak cross-correlation z-scored against GLM simulation).

In formulas, denoting with i the mPFC cell for which the prediction has to be made, and with $\lambda_1, \dots, \lambda_N$ the firing rate of N functionally coupled CA1 rate maps, and w_{ij} their functional correlation strength, then we defined the predicted mPFC map as

$$\lambda_i^{pred} = \sum_{j=1}^N w_{ij} \lambda_j$$

We then reported the Pearson correlation between real λ_i and predicted firing rate map λ_i^{pred} .

Spatial information measure

We computed the spatial information per spike.^{64,65} In a nutshell, this corresponds to the first order approximation of the mutual information between position-dependent spiking probability and location, divided by the average firing rate. Denoting with $x \in \mathcal{X}$ the location, and $\lambda(x)$ the average firing of a cell at location x , and with $\lambda = \langle \lambda(x) \rangle_x$ its average firing rate across locations, the spatial information is defined as

$$\left\langle \frac{\lambda(x)}{\lambda} \log \frac{\lambda(x)}{\lambda} \right\rangle_x$$

where $\langle \cdot \rangle_x$ denotes an average over positions.

CA1 theta oscillations

We averaged the downsampled (5KHz) LFP signal for each tetrode that was implanted in CA1. We then extracted the theta signal (5–15 Hz) by means of a 3rd-order Butterworth filter. We detected the local minima of the theta oscillation, and linearly interpolated between each consecutive pair of local minima from 0 to 360°.

Theta locking strength

We computed the theta (θ) phase of each spike by detecting the angle (0–360) of the underlying theta oscillation at the time of activity. We computed phase histograms by binning angles into 30° bins. We quantified theta locking strength by means of the mean vector length (MVL): this is defined as the absolute value of the average angle (appropriately converted in radians) in the complex plane:

$$MVL = \text{abs} \left(\frac{1}{N} \sum_{x=1}^K \exp(i \theta_x) \right)$$

Theta locking and optimal delay

We employed a measure introduced by Siapas and colleagues in²⁸ to obtain the optimal delay of mPFC cells theta-locking relative to hippocampal theta. We employed a Rayleigh test for circular uniformity and selected all mPFC cells that yielded a < 0.05 . For each one of those cells, we introduced a delay $\tau \in \{-125, -124, \dots, 124, 125\}$ ms to all spikes, again computed the phase as described above and computed the MVL of these angles. We then selected the delay that yielded the highest MVL for each cell.

Z-scored theta locking strength of paired vs. independent spikes

Consider a mPFC-CA1 functionally correlated cell pair. We defined a mPFC spike as being “paired” if it happened in a theta cycle when also the functionally correlated CA1 emitted at least one spike.

We tested whether paired spikes had better theta phase locking than independent spikes. To do so, we measured the MVL of the theta-phases of paired spikes, and z-scored it against 200 MVLs measured on randomly subsampled independent spikes. In the (rare) case where paired spikes were more than independent ones, we randomly subsampled paired spikes instead, and reported the Z score with sign flipped. Excluding these rare cases did not alter the results.

Theta-dependent functional correlation

We employed the same method described above, with the difference that we considered across area cross-correlations only for CA1 spikes that happened at a particular theta phase. In detail, we binned angles into 12 equally sized bins (36°). For each bin $[\theta_m, \theta_M]$, we set to zero all the CA1 activity that happened outside $[\theta_m, \theta_M]$. We then proceeded with the quantification of functional correlations as described above.

Average theta distance between pairs of cells

For each pair of functionally correlated CA1-mPFC cells, we measured the difference in theta phase among each pair of spikes that happened within the same theta cycles. We then measured the average angular difference as a function of CA1 phase, and reported angular standard error of the mean in the figure.

The difference between radian angles θ_1 and θ_2 is defined as:

$$\theta_1 - \theta_2 = \arg(\exp(i\theta_1) - \exp(i\theta_2))$$

The average of a set of k radian angles $\theta_1, \dots, \theta_k$

$$\bar{\theta} = \arg\left(\sum_{j=1}^k \exp(i\theta_j)\right)$$

The circular variance (for radian angles) is defined as $\text{Var}(\theta_1, \dots, \theta_N) = 1 - \text{MVL}(\theta_1, \dots, \theta_N)$, and the standard deviation of the circular mean as $\sqrt{-2 \log \text{MVL}(\theta_1, \dots, \theta_N)}$. In the figures, we reported all angles in degrees (0.360), transformed from radians to degrees by using: $\frac{(\theta + \pi)}{\pi} \cdot 360$

Randomization of theta phase and drop of correlations

To test whether precise theta timing was required for mPFC cells to be functionally correlated with CA1 cells, we randomized the theta locking. To do so, for each mPFC cell independently we considered different levels of random spike jittering. In detail, to each spike we added a random time uniformly distributed within $[-M, M]$, where M is the maximum jitter allowed. Afterward, we computed CA1-mPFC functional correlations using mPFC-jittered spike trains, while leaving CA1 untouched. We then measured the reduction of functional correlation strength by considering the relative decrease in magnitude of the peak: $\frac{\text{peak}_{\text{orig}} - \text{peak}_{\text{jitter}}}{\text{peak}_{\text{orig}}}$

Within theta-cycle shuffling

We employed a shuffling method that preserves the measured theta-phase histogram of each cell.

For each cell, during each theta cycle, we used Inverse transform sampling (based on the empirical theta-phase histogram) to sample the same number of spikes that were originally fired.

This shuffling procedure is very conservative in that it preserves every aspect of neuronal firing (e.g., spatial, speed, head direction, etc.) that is slower than theta rhythm (5–15 Hz).

Power spectral density

We employed Welch’s method to compute averaged power spectral density.⁶⁶ In a nutshell, this method works by dividing the signal into overlapping segments and averaging the periodograms obtained in each. In detail, we performed short-time Fourier transform centered at each theta cycle, with a Hann window of 400ms, separately for all CA1 and PFC tetrodes. For each CA1-mPFC cell pair, we computed the power spectrum for theta cycles where both cells are active (paired) vs. theta cycles where either or both cells are silent (independent). Finally, we averaged the periodograms across cells, separately for CA1 and mPFC.

Sensitivity of cells to gamma and sharp-wave ripple oscillations

We measured the average firing rate of mPFC cells during periods of when gamma (40~120Hz) or SWR oscillations (150–250 Hz) had high power (>95th percentile) and compared to the average firing rate. Gamma and SWR power was measured as the average power spectrum computed from short-time Fourier transform of the CA1 LFP during time windows of 500 ms and Tukey windows with shape parameter of 0.25.

Bipartite spectral clustering

We considered a graph where each node represents a cell, and edges are present only among CA1-mPFC cell pairs that were functionally correlated. We then asked if it was possible to subdivide this (bipartite) graph into smaller clusters, containing both CA1 and mPFC cells, so as to maximize the across-area connections within each cluster and minimize the ones outside. We employed a method based on the spectral optimization of modularity (SOM) algorithm, first proposed in.³² The algorithm iteratively finds the best split by considering the eigenspectrum of the modularity matrix M of the bipartite graph. The extension of this work to bipartite graphs was provided in.⁶⁷ Intuitively, the modularity of a graph measures the tendency of vertices to cluster into modules, or, in our case, neurons into assemblies. The eigenspectrum of B holds properties that allow for an easy implementation of an iterative algorithm for identifying modules in bipartite networks.

Details: consider a graph with n vertices and m edges, and adjacency matrix A . Each vertex is assigned to a module g_i . The measure that one aims to minimize is the modularity Q , defined as

$$Q = \frac{1}{2m} \sum_{ij} (A_{ij} - P_{ij}) \delta(g_i, g_j)$$

where $\delta(g_i, g_j) = 1$ if and only if $i = j$, and zero otherwise. P_{ij} represents the probability of a null model of finding an edge between i and j .

An alternative definition of modularity, in matrix form, can be defined as follows

$$M = \frac{1}{2m} \text{Tr}(S^T B S)$$

where $B_{ij} = A_{ij} - P_{ij}$ and $S = [s_1, \dots, s_c]$, for modules, is an n -by- c matrix determining the belonging of each node to each cluster. Barber (2008) derived the proper choice of P_{ij} for the bipartite cluster, which is $P_{ij} = \frac{k_i h_j}{m}$, where k_i, h_j represent the number of incoming edges to nodes i and j , respectively. The algorithm then iteratively finds the best split of the graph into two subgraphs by finding the largest positive eigenvalue of M and the corresponding eigenvector v ; the sign of each entry i of v determine the belonging of node i to the first or the second module. The iterative algorithm stops after the modularity measure stops growing, or when modules sizes get too small (<3), or when the number of iterations goes beyond a certain number (>6).

Cleanliness of assemblies

For each assembly, we measured the average proportion of connections within assembly (0–1) and subtracted the average proportion of connections outside the assembly (0–1). The resulting measure can range from -1 (very bad) to 0 (same connectivity within vs. outside) to 1 (perfect connectivity within and no connectivity across assemblies).

Assembly firing rate map

We deemed an assembly active whenever at least 2 CA1 and 2 mPFC cells were active in the same theta cycle. Afterward, we followed the same prescription described above for computing firing rate maps, with the difference that instead of using spike numbers we considered only 1 (active) or 0 (not active) during a given theta time window.

Within area functional correlations

We employed the same statistical procedure presented above for within-area functional correlations.

Bootstrapped theta locking difference

For each pair of mPFC cells belonging to the same assembly, we randomly selected 50 spikes from each cell and measured their average phase. We did so 200 times for each cell pair. Afterward, we counted how many times the average phase of the first cell was later than the average phase of the second cell (i.e., the difference of the two cell was $<180^\circ$). If the two cells had no phase relationship among each other, the probability of one phase being larger than the other would be $p = 0.5$. If that were the case, 99% of the cases one would have a count between ~ 75 and ~ 125 . We quantified the percentage of cell pairs that were outside this confidence interval, and reported the Binomial test with $p = 0.01$ baseline.

Activity of assemblies beyond expected by GLM model

We summed the co-activity of all CA1-PFC pairs of cells belonging to the same assembly in consecutive 25ms time windows and compared this co-activity against what we expected from 1000 GLM simulations. We deem an assembly significantly active at a time window when its coactivity was larger than $\text{MEAN} + 3\text{STD}$ of the simulated GLM co-activity. Using this result we calculated the proportion of (significantly) active assemblies for each time window, which was then summarized in 5 macro-regions of the linearized arena (0–40cm, 40–80cm, 80–120cm, 120–160cm, 160–200cm). In Figures 7A and 7B, we compared the average co-activity during correct and error trials. In Figures 7D and 7E, we compared the average co-activity during the one trial preceding, corresponding to, and succeeding a strategy change. The trial of strategy change was defined as the first trial when the animal stopped following the previous rule and kept a consistent behavior afterward (e.g., stop following the light and consistently choose to go west, Figure 7C).

Modeling

GLM model of cells responses

We utilized a GLM model to describe each cell's response propensity as a function of all measured covariates during foraging activity. These detailed models served as null models for the statistical test we employed to detect functional correlations among cells (see

section “Detection of across-area couplings”). We will describe in detail here the covariates used to fit the model and the parameters used in fitting and simulation routines, and refer the reader to other references for the details regarding the theoretical background of GLMs.⁶²

The model described the inhomogeneous Poisson activation rate λ_t of cells in 25.6ms time windows

$$P(x_t = n | \lambda_t) = e^{-\lambda_t} \frac{\lambda_t^n}{n!}$$

The expected firing rate λ_t takes the form $\lambda_t = \exp(\beta \cdot \chi_t)$ where β represents the model coefficients, found by maximum likelihood (below) and χ_t represents the vector of covariates, which is a concatenation of the following:

- trajectory-dependent spatial position: we binned the linearized position in 10cm bins, as described above, in 25.6ms time bins. We allowed each cell to have a different encoding for each trajectory separately to allow for maximum flexibility: this resulted in a 80-dim vector, i.e., 20 location bins for each possible trajectory. At each time point, only the entry of this 80-dim vector that represented the location of the animal, and the trajectory taken, was set to 1, and all the others to 0 (one-hot encoding variable)
- speed: we binned the speed, which was measured from the behavioral recording in 25.6ms time bins, in 7 non-overlapping and equally populated speed bins, starting from 7 cm/s (one hot variable)
- theta phase: we computed the theta phase at the center of each 25.6ms time bin. The computation of the theta phase is detailed in a previous section. We binned the angles in 10 non-overlapping angular bins, and encoded it as a one-hot variable.
- spiking history: the spiking of the last three time bins was used.
- within-area spiking activity: the spiking activity of all the other cells in the same area, together with the spiking history of each of those cells in the previous time bins were used as additional covariates.

The number of parameters of such models ranged from 100 to 400, depending on the number of cells recorded simultaneously in the same area. We utilized the routine GLM offered by the package statsmodels v0.12.2 in Python 3.7. We fitted the models by using an L2 regularization, whose parameter was found by grid search on $(10^{-10}, 10^{-9}, \dots, 10^0)$ and cross-validation (train = 75%, test = 25% of data) independently for each cell.

Deviance explained

Using the same notation as above for the expected firing rate λ_t , and denoting with y_t the actual number of spikes observed, the deviance of a Poisson GLM is defined as

$$D = 2 \sum_{t=1}^T \left[y_t \log \frac{y_t}{\lambda_t} - (y_t - \lambda_t) \right]$$

This measure captures the deviation of the data from a Poisson GLM model⁶⁸ (McCullagh and Nelder 1989). This measure can be compared to the deviance explained by a simple null model that always predicts the mean firing, denoted here by $\bar{\lambda}$:

$$\bar{D} = 2 \sum_{t=1}^T \left[y_t \log \frac{y_t}{\bar{\lambda}} - (y_t - \bar{\lambda}) \right]$$

The deviance explained by the full GLM model is the defined as

$$ED = 1 - \frac{D}{\bar{D}}$$

Finally, to measure the relative contribution of the various covariates (e.g., position, speed) we divided the ED of the partial model by the ED of the full model.

Normative model of (spatial) information transfer and generalization

Consider two populations of neurons, exemplifying N PFC cells and M CA1 hippocampal cells. We will denote with $x = (x_1, \dots, x_N)$ the stochastic (binary) activation of PFC cells, and with $y = (y_1, \dots, y_M)$ the stochastic (binary) activity of CA1 cells. Let us denote with s the distance to the goal, and with k the trajectory. We consider s to be a random variable that takes values in $\{1, \dots, 10\}$, each with equal probability, and k a uniform Bernoulli $\in \{0, 1\}$ independent of s .

We assume that x is driven by hippocampal inputs and internal connectivity. We formalize this with a stochastic model, which is similar in its formalization to a restricted Boltzmann machine:⁶⁹

$$P(x | y) \propto \exp \left(\sum_{i=1}^N \sum_{h=1}^M C_{ih} x_i y_h + \sum_{i,j=1}^N W_{ij} x_i x_j \right)$$

Where $C \in R^{N \times M}$ is the feedforward hippocampus to mPFC connectivity, and $W \in R^{N \times N}$ the internal mPFC recurrent connectivity.

We assume that the activity of CA1 cells is both trajectory and position dependent, and we find the trajectory-independent firing by marginalization: $P(y | s) = \sum_{k=0}^1 P(y | s, k) P(k)$

The mPFC population activity x inherits position dependency through activity y , i.e.,: $P(x | s) = \sum_y P(x | y) P(y | s)$

For fixed W and given $P(y | s, k)$, we want to find the best C that maximizes the mutual information between x and s : $MI(x, s) = \sum_{x, s} P(x, s) \log \frac{P(x, s)}{P(x)P(s)} \propto \sum_{x, s} P(x | s) \log \left(\frac{P(x | s)}{P(x)} \right)$

We maximize this quantity via Sequential Least Squares Programming (SLSQP) routine in SciPy. We constrain each C_{ij} to lay in $[-1, 1]$. Our simulations use $N = M = 10$ neurons, which allows the mutual information to be computed without the need for approximations (by enumerating all possible patterns).

Reported estimates are obtained by averaging across 100 randomly initialized networks; for each simulation, $\{C_{ij}\}$ is a symmetric matrix with entries randomly sampled from a $\{0, 1\}$ distribution, and $\{C_{ij}\} P \left(\begin{matrix} \rightarrow \\ y \end{matrix} \middle| s, k \right)$ Gaussian place field per trajectory and simulated in such a way to resemble CA1 single-cell statistics measured in the data.

Programming languages and packages used

All scripts were written in Python 3.7. We used the packages NumPy, SciPy, StatsModels, and AstroPy.