

Prefrontal Cortical Neurons Are Selective for Non-Local Hippocampal Representations during Replay and Behavior

Alice Berners-Lee,^{1,2} Xiaojing Wu,² and David J. Foster¹

¹Helen Wills Neuroscience Institute and Department of Psychology, University of California, Berkeley, California 94720, and ²Solomon H Snyder Department of Neuroscience, Johns Hopkins University School of Medicine, Baltimore, Maryland 21205

Diverse functions such as decision-making and memory consolidation may depend on communication between neurons in the hippocampus (HP) and prefrontal cortex (PFC). HP replay is a candidate mechanism to facilitate this communication, however details remain largely unknown because of the technical challenges of recording sufficient numbers of HP neurons for replay while also recording PFC neurons. Here, we implanted male rats with 40-tetrode drives, split between HP and PFC, during learning of a Y-maze spatial memory task. Surprisingly, we found that in contrast to their non-selectivity for maze arm during movement, a portion of PFC neurons were highly selective for HP replay of different arms. Moreover, PFC neurons' selectivity to HP non-local arm representation during running tended to match their replay arm selectivity and was predictive of future choice. Thus, PFC activity that is tuned to HP activity is best explained by non-local HP position representations rather than HP representation of actual position, providing a new potential mechanism of HP-PFC coordination during HP replay.

Key words: decision-making; hippocampus; PFC; reactivation; replay

Significance Statement

The hippocampus (HP) is implicated in spatial learning while the prefrontal cortex (PFC) is implicated in decision-making. The question of how the two areas interact has been of great interest. A specific activity type in HP called replay is particularly interesting because it resembles internal exploration of non-local experiences, but is technically challenging to study, requiring recordings from large numbers of HP neurons simultaneously. Here, we combined replay recordings from HP with prefrontal recordings, to reveal a surprising degree of selectivity for replay, and a pattern of coordination that supports some conceptions of hippocampocortical interaction and challenges others.

Introduction

Neurons in the hippocampus (HP) represent an animal's location while it traverses space (O'Keefe, 1976). During rest periods, sharp-wave ripples (SWRs; 150–250 Hz) in the HP local field potential (LFP) occur concurrently with patterns across populations of HP neurons that depict trajectories through explored space, termed "replay" (Nádasdy et al., 1999; Foster and Wilson, 2006; Karlsson and Frank, 2009; Foster, 2017). Replay trajectories are non-local, and can explore space in ways the animal never did (Foster and Wilson, 2006; Gupta et al., 2010). Additionally, SWR disruption causes spatial memory deficits (Ego-Stengel and Wilson, 2010; Jadhav et al., 2012). Therefore, SWRs and

associated replay have been suggested to support complex processes such as memory-guided decision-making (Singer et al., 2013) and systems memory consolidation (Buzsáki, 1996; Sutherland and McNaughton, 2000; Teyler and Rudy, 2007). However, mechanistic insights into how HP replay supports these functions has been lacking.

SWRs are substantial spiking events that broadly engage neocortical areas (Logothetis et al., 2012). One of these, the prefrontal cortex (PFC), has been implicated in memory consolidation and decision-making (Miller and Cohen, 2001; Vertes, 2006). Individual PFC neurons are modulated by SWRs (Jadhav et al., 2016; Tang et al., 2017), and more so during tasks than in subsequent rest, and more during initial phases of learning (Jadhav et al., 2016; Tang et al., 2017). Furthermore, PFC neurons that are engaged in a task and locked to HP theta are more likely to be activated by SWRs (Jadhav et al., 2016). PFC neurons are modulated more to SWRs co-occurring with the reactivation of movement-related places than immobility-associated places (Yu et al., 2017). Similarly, coherence of co-activity between the two areas during SWRs tends to be stronger for SWRs that reflect the path the animal will take or has taken than those that reflect

Received May 12, 2020; revised Apr. 24, 2021; accepted Apr. 26, 2021.

Author contributions: D.J.F. designed research; X.W. performed research; A.B.-L. analyzed data; D.J.F. and A.B.-L. wrote the paper.

This work was supported by the National Institutes of Health Grant R01MH085823.

X. Wu's present address: New York University Comprehensive Epilepsy Center, New York, New York 10016.

The authors declare no competing financial interests.

Correspondence should be addressed to David J. Foster at davidfoster@berkeley.edu.

<https://doi.org/10.1523/JNEUROSCI.1158-20.2021>

Copyright © 2021 the authors

a not-taken path (Shin et al., 2019). These studies suggest that both awake replay and PFC activity are involved in processing place information related to the current task. Further, HP-PFC neuron pairs that have more correlated firing patterns across an entire environment are more likely to fire together during SWRs (Peyrache et al., 2009; Jadhav et al., 2016). However, without the ability to separate the representation of different places within SWRs, it remains unclear whether PFC neurons are selective for places represented in the HP population, and under what conditions. We hypothesized that by deploying sufficient numbers of tetrodes into HP to resolve the place content of HP replay, in a Y-maze environment already shown to produce clearly segmented replay of different “arms” of the track, we could leverage the robust place representation in HP to identify the relationship of concurrently recorded PFC units to HP replay information content.

Here, we recorded from neurons in prelimbic and infralimbic PFC areas in rats on novel tracks during early learning, while simultaneously recording enough HP neurons to resolve trajectories depicted during HP replay events and HP theta sequences. We found that although PFC neurons were relatively non-selective for spatial position during behavior, they were surprisingly selective for different parts of the environment during HP replay. These data are the first to our knowledge to find that individual PFC neurons can be selective to replay of different spatial trajectories within an environment (here, arms in a Y-maze). Furthermore, during active behavior, when the HP was representing non-local information, individual PFC neurons exhibited the same spatial selectivity as during HP replay. Thus, PFC neurons’ lack of spatial selectivity was a consequence of selectivity to non-local place representation in HP, which is, by definition, unrelated to the current location of the animal. These findings support the idea that PFC neurons’ responses to HP replay are important for ongoing behavior in early learning. Moreover, they introduce the notion of consistent tuning of PFC responses to internally generated, non-local representations, with ramifications for our understanding of the function of HP-PFC interactions.

Materials and Methods

Data from four of the 11 sessions analyzed here were also used in a previous study (Wu and Foster, 2014). The recording and preprocessing methods in this paper are identical to that study and are re-stated here.

Experimental design and statistical analysis

Four male Long-Evans rats, 3.5–4.5 months in age, were pretrained under moderate food deprivation to run back and forth on a linear track for chocolate milk reward available at both track ends. Custom-designed micro-drives consisting of 40 independently adjustable tetrodes were then implanted with half the tetrodes targeting the right dorsal CA1 area of the HP (−3.6, −2.2 mm from bregma), and the other half targeting the right medial PFC (+3.2, −0.8 mm from bregma). Tetrodes were gradually moved into the CA1 pyramidal cell layer with one tetrode left in the white matter above the cell layer to provide a reference signal to all other tetrodes. Once all tetrodes were in position, rats were exposed to a modified Y-maze for the first time, in a novel room with distal visual cues in the form of shapes on the walls and equipment around the room, and recording began. The Y-maze was composed of one long arm (161 cm) and two short arms (81 cm), all of which were separated by 120°. One short arm was chosen to be the central (C) arm; the other two arms [right (R); left (L)] were termed “alternating arms.” The rat was placed at the baited end of the central arm, and was allowed to freely explore the Y-maze. The rat was rewarded according to the following alternation rule: the first arrival at an alternating arm was rewarded; thereafter, returns to the central arm were rewarded while visits to the

alternating arms were rewarded only if the arm identity was different from the last alternating arm rewarded. Alternation accuracy is defined as the percent of correct trials out of all trials where the rats left the center arm (and chose one of the two alternating arms). Spiking activities and LFP signals were recorded (Neuralynx) and digitally filtered in different frequency bands (600–6000 Hz for spikes and 0.1–500 Hz for LFP). Rats’ positions were signaled by red and green LEDs attached to the drive, and were recorded from an overhead camera. Recording sessions occurred each day and lasted 30 min to 2 h. There were no specific criteria for ending each session. The rats explored all three arms in each day and performed an average of 68.7 passes (reward to reward) on the track each day and an average of 6.6 correct trials (center, right, center, left). Recordings were terminated when HP neuron yield became poor after 6–8 d. Eleven sessions across four rats were used (one to four sessions per rat). Rats were killed, and lesions were made on selected tetrodes by passing current (30 A for 5 s) through each tetrode. Animals were then perfused with formalin. Brains were removed, sectioned, and stained for cresyl violet to verify tetrode recording locations. The spike data were manually clustered in custom software Xclust2 (written by M. Wilson, Massachusetts Institute of Technology, Cambridge, MA). Putative HP pyramidal neurons were separated from HP interneurons by their spike width and spike rate. None of the HP interneurons were considered in this analysis. Well isolated single units recorded in the PFC were included in subsequent analyses if the following criteria were met: (1) mean peak-to-trough spike width > 0.35 ms; (2) fewer than 5% of all spikes occurred within the refractory period (interspike interval < 2 ms); (3) average firing rate over the entire recording session is higher than 0.2 Hz. These criteria were set to ensure exclusion of noisy clusters, units with ultralow firing activities, and interneurons. Eleven sessions across four rats were used, except where otherwise noted where only the session with the highest yield of PFC neurons was used from each rat (Fig 3H, inset, Fig. 4F, inset, Fig. 6K,M, insets).

All analysis and statistics were performed in MATLAB and are described in the following sections. Data and MATLAB code are available from the authors on request. All permutation tests were done by taking the proportion of permutations (n) that produced a value (perm) greater than or equal to the real value the data produced (real); $P = (\text{sum}(\text{perm}) \geq \text{real}) / (n + 1)$. All error bars and shading represent SEM.

Position linearization and place-field computation

Recorded positions were projected onto three centerlines (defined by the experimenter) aligned with the three arms of the Y-maze. The three lines were then concatenated to produce a linear axis (Fig. 3A). Directional linear place fields were computed by using all spike and position data (2-cm position bins) from when the rat was moving >5 cm/s and smoothed with a Gaussian filter ($\sigma = 3$ bins). Note that firing rates on the three arms were separately smoothed to minimize misestimation of place fields around the choice point. Two-dimensional place fields were computed by using all spike and position data (8×8 cm position bins) from when the rat was moving >5 cm/s and smoothed with a Gaussian filter ($\sigma = 2$ bins; Fig. 1B,C).

Spatial selectivity and radial symmetry

Spatial selectivity was calculated for each neuron by taking the maximum value across linear position bins divided by the mean value across those same position bins (Fig. 1D). Radial symmetry was calculated for each neuron by taking the average Pearson’s correlation value across the linear firing rate maps for all pairs of the three arms (the long arm was binned by an additional factor of two for comparison; Fig. 1E). Comparisons of these metrics between HP and PFC populations were done using two-sided Wilcoxon rank-sum tests.

Overlap measure

In addition to radial symmetry, we evaluated the extent of overlap between fields across arms using the linearized firing fields and a method previously published by Battaglia et al. (2004) to produce an overlap (r) evaluated with no shift ($s = 0$). We took the average overlap across the three combinations of arms (e.g., center compared with right arm, again

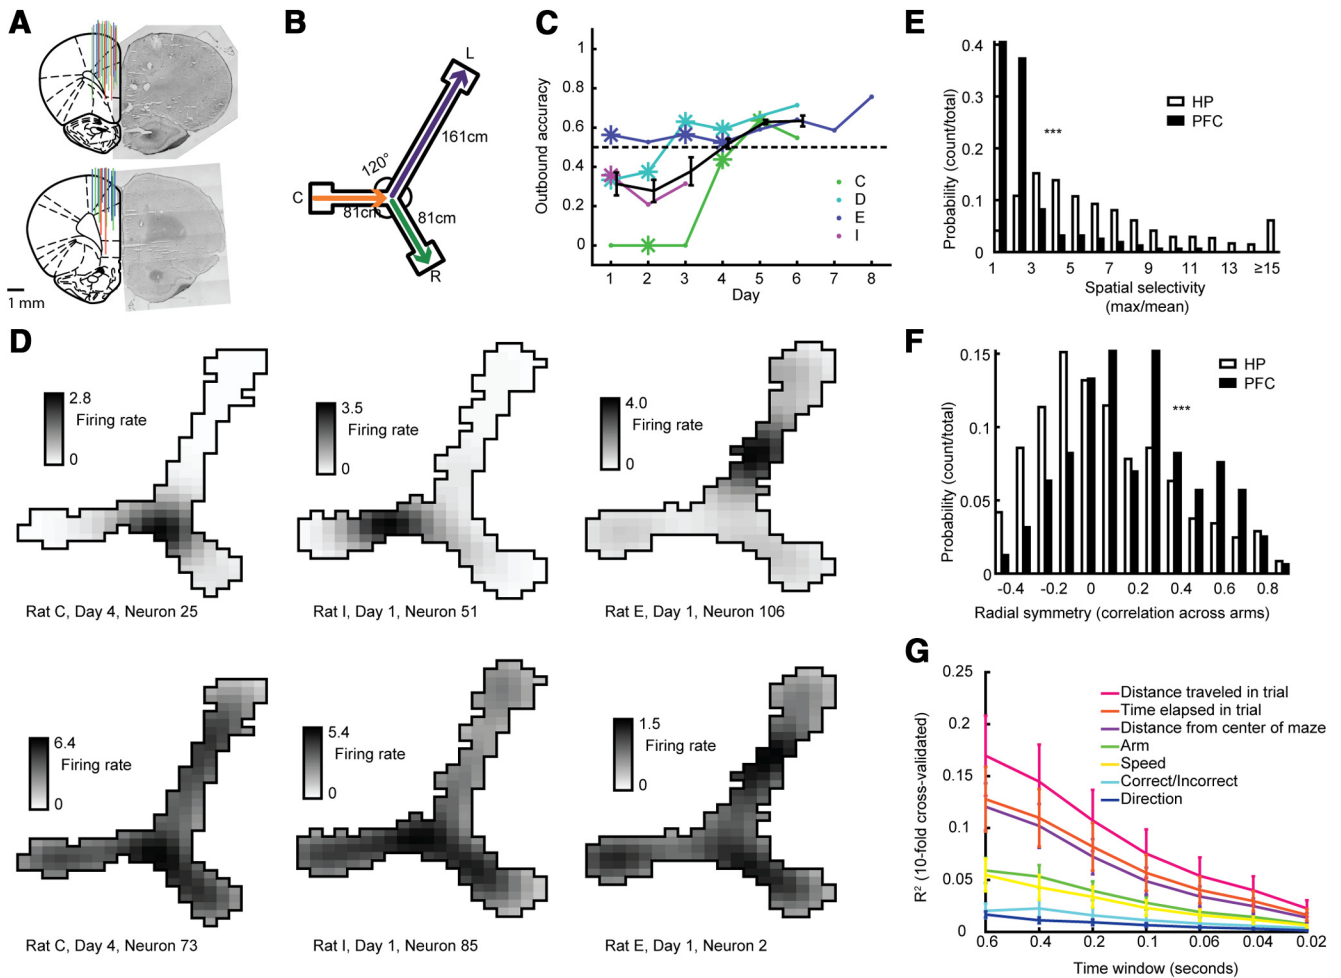


Figure 1. HP neurons were more modulated by the rat's position on the track than PFC neurons. **A**, PFC tetrad placement. Left, Summary of three rats, each with different color. Right, example histology images. **B**, Rats were rewarded for alternating between right and left arms, and always rewarded for returning to the center arm. **C**, Outbound accuracy in alternation task across days. Asterisks depict the 11 d from which neural recordings were used. Kendall's rank correlation between recording day and outbound accuracy, $R = 0.56$, $p = 3.6 \times 10^{-4}$. **D**, top, Three examples of HP neurons recorded from three different rats showing place fields. Bottom, Three example prefrontal (PFC) neurons showing little spatial selectivity recorded simultaneously to the HP neurons in **B**. **E**, Spatial selectivity (max firing rate in binned position divided by the mean firing rate across bins) for all HP and PFC neurons. Two-sided Wilcoxon rank-sum test, $N = 158, 968$, $z = 5.36$, $p < 10^{-7}$. **F**, Radial symmetry (correlation across the three arms) for all HP and PFC neurons. Two-sided Wilcoxon rank-sum test, $N = 158, 968$, $z = -14.35$, $p < 10^{-45}$. **G**, Results from 10-fold cross-validated predictions of behavioral correlates using PFC neurons' firing rates.

binning the long arm by twice as much to match the shorter ones) for each PFC cell to generate the overlap measure.

Generalized linear model (GLM)

To explore what behavioral aspect was best predicted by PFC neuron firing we performed a GLM using all PFC neurons' firing rate each lap (from when rat left reward area to when the rat reached the next one) binned into time bins 0.02–0.6 s long. A GLM was created from this activity to predict one of seven behavioral variables. The MATLAB function *fitglm* was used and the three categorical variables were marked as such in the function. We used a 10-fold cross-validation whereby 90% of the data were used to predict the remaining 10%. The square of the correlation value of the true data and the data predicted by the GLM was plotted in Figure 1G.

Candidate event and replay analysis

For each recording session, we identified the three tetrodes from which the most HP neurons were isolated. The LFP from these tetrodes was bandpass filtered between 150 and 250 Hz, and the absolute value of the Hilbert transform of this filtered signal was then smoothed (Gaussian kernel, $SD = 12.5$ ms). To examine SWRs, these processed signals were averaged across all three tetrodes and SWRs were identified as local peaks with an amplitude > 2 SD above the mean, excluding periods

when the rat's speed was > 5 cm/s. The start and end boundaries for each event were defined as the point when the signal crossed the mean. SWRs shorter than 50 ms or longer than 2 s were excluded from further analysis. A Bayesian decoding algorithm (Davidson et al., 2009) was then applied to the candidate events, which calculated the probability of the ensemble of neurons representing each position bin during each time bin (20-ms time bin, every 5 ms).

To identify SWRs that represented a certain arm (termed "arm-replays"), we segmented the posterior probability matrix of each candidate event in both position and time to further define trajectory-specific subregions. Each SWR event was first separated along the position axis into three segments corresponding to the three arms. We defined a maximum a priori probability (MAP) function as the largest probability across all positions per time bin, smoothed in time with a Gaussian (Wu and Foster, 2014). Each single-arm segment was then segmented in time, and a trajectory-specific subregion was created around the largest peak of the MAP function to include time bins whose MAP values were above a threshold of four times the chance level, namely $(1/\text{total \# position bins}) > 4$ (a fixed threshold only dependent on the size of the track and not the quality of neuronal data). To identify joint-replays, for all arms that sustained at least 50 ms above the MAP threshold the following three variables were calculated: (1) length of subregion in time (> 50 ms); (2) arm coverage ($> 50\%$ of the arm); and (3) weighted correlation

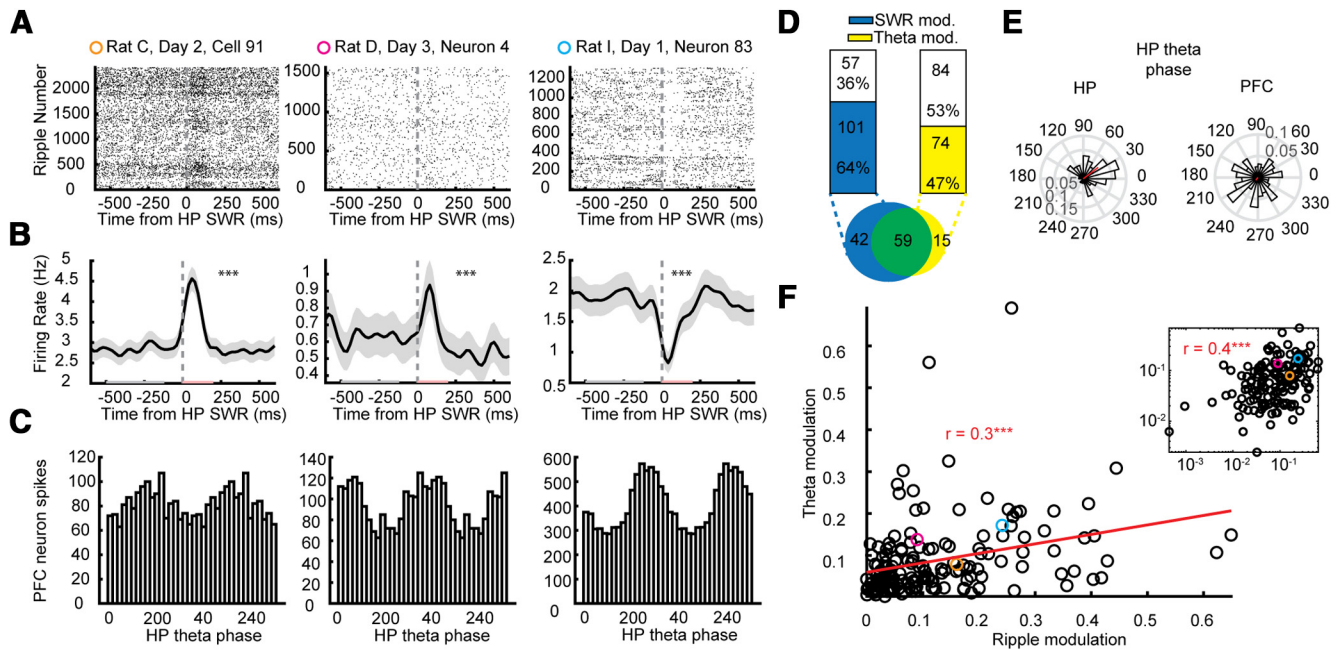


Figure 2. PFC neurons were modulated to HP ripple events and HP theta. **A**, Three example PFC neurons' raster plot of firing centered on HP ripple events. **B**, Average ripple-triggered firing rates of the same three PFC neurons in **A**; p values are from a permutation test using a distribution of circularly shifted firing rates. **C**, Histograms of the same three PFC neurons' spikes as a function of HP theta phase. **D**, left, 64% of PFC neurons were significantly modulated by HP ripples (permutation test $p < 0.05$); 58% of ripple-modulated neurons were modulated by theta compared with only 26% of non-ripple-modulated neurons. Right, 47% of PFC neurons were significantly modulated by HP theta (Rayleigh test $p < 0.05$); 80% of theta-modulated neurons were modulated by ripples compared with only 50% of non-theta-modulated neurons. Inset below, Venn diagram of neurons modulated by ripples (blue), theta (yellow), and both (green). **E**, Histogram of the mean theta preference for each HP neuron (top) and PFC neuron (bottom) to HP theta. **F**, The more a PFC neuron was modulated to HP ripples (absolute value of the FR in postripple (red in **B**) minus preripple (gray in **B**); period normalized by sum of both) the more it tended to be modulated to HP theta (mean resultant vector). Pearson's correlation, $R = 0.30$, $p = 1.3e-4$. Inset, Same data plotted on a log-log scale (Person's correlation: $R = 0.42$, $p = 6.2e-8$). The three examples shown in the left, middle and right panels of **A–C** are depicted in orange, pink, and blue, respectively.

(absolute value > 0.3). If two non-overlapping replays satisfied these conditions and the resulting joint-replay's absolute value of its weighted correlation exceeded 0.3 then it was deemed a joint-replay. These replays were separated and for all future analyses as two separate single-arm-replays, each identified as an arm-replay. All SWRs that did not have joint-replays but had one arm that stayed above the MAP threshold for at least 50 ms were truncated in time to when it crossed the MAP threshold and identified as an arm-replay. If no arm crossed the MAP threshold for in a SWR it was not categorized as an arm-replay. We assessed whether PFC neurons had a bias in modulated to arm-replays depending on their direction of travel, their dependence on a certain directional map, whether they were forward or reverse replays or whether they were arm-replays going toward or away from the rat on the arm it was on and found no bias in any of these comparisons. Because of this, all arm-replays, despite the direction of travel of the replay, the directional map it used, or whether it was using the same (forward replay) or different (reverse replay) map as the direction of travel, were used.

PFC SWR modulation

PFC neurons were categorized as being significantly SWR modulated as was done by Jadhav and colleagues (Jadhav et al., 2016), briefly restated here. We averaged over all SWRs to create peri-SWR time histogram (SWR-PSTH). To test significance, we created 500 shuffled SWR-PSTHs. In each SWR, the PFC spikes were jittered circularly around the SWR (in the window of -2 – 2 s around an event) by a random amount. We calculated the squared difference between the real SWR-PSTH and the mean of the shuffled SWR-PSTHs in the 0- to 200-ms window after SWR onset, to obtain a SWR-modulation measure. We then repeated this for each differently jittered SWR, producing 500 shuffled modulation values. If the real SWR modulation exceeded 95% of the shuffled distribution (permutation test), then the neuron was significantly modulated (Fig. 2A,B,D). The direction of that modulation was determined by the sign of the SWR modulation. To compare the SWR modulation to theta modulation across neurons we computed the difference between

the 200 ms after SWR onset to a baseline period in the -500 - to -100 -ms window before the SWR and normalized it by dividing by the sum of those two periods (Fig. 2G).

PFC arm-replay modulation

To assess whether individual PFC neurons were differentially modulated to SWRs depicting different arms, arm-replay-PSTHs were made for each neuron. Using the same time windows as with SWR modulation, the modulation of each PFC neuron to each of the three arm SWRs (SWRs depicting the three arms of the Y-maze) was computed and averaged across each arm-replay category. The modulation was defined as the difference in firing rate in a modulation window (200 ms after replay onset) from a baseline period (-500 to -100 ms before the replay). To assess the significance of the difference in modulation, the sum of squared differences between all pairs of the three arms was taken, and this was summed to produce a single value (SSD). This SSD was then compared with 2000 shuffled distributions of permuting the group label (which arm is depicted) of each arm-replay event and repeating the SSD analysis. A PFC neuron was determined to be significantly differentially modulated by arm-replays if the true SSD exceeded 95% of the distribution of shuffled SSDs (permutation test; Fig. 3). Neurons that did not have sufficient spikes (< 50) summed across all SWRs in both windows combined were not included in this analysis. We also replicated these analyses using only arm-replays that passed the stricter criterion of a weighted correlation absolute value > 0.6 and max jump distance < 0.4 and found similar and significant results with 22.7% of PFC neurons significantly differentially modulated by replays of different arms.

Expected PFC arm-replay modulation assuming re-activation

To simulate the expected modulation of PFC neurons to arm-replays as re-activation of arm-selectivity during running, we kept the proportion of arm-replays the same, but shuffled the labels of them non-randomly to minimize the difference of the pattern of arm-replay modulation elicited and the pattern the PFC neurons exhibited across running.

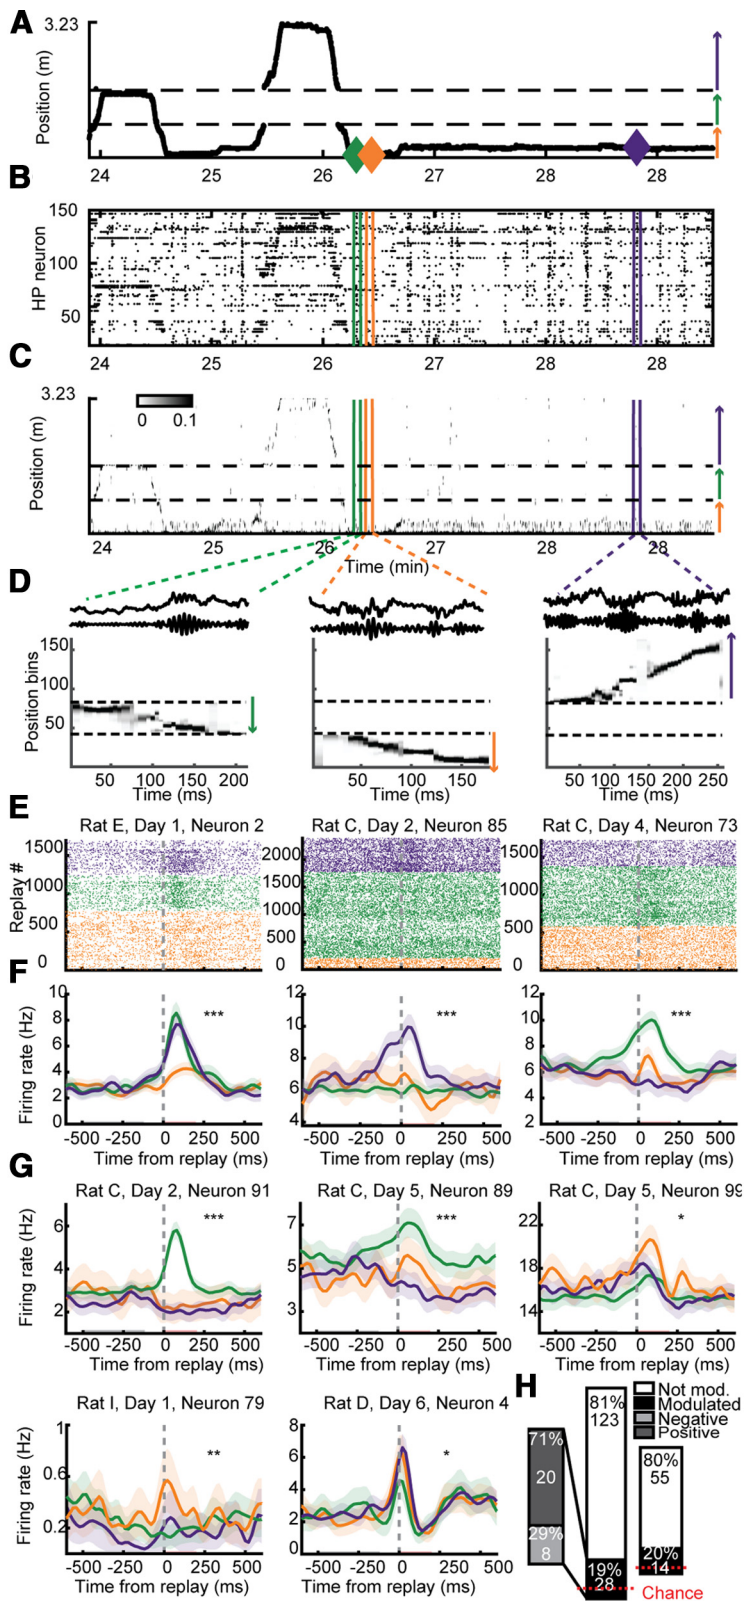


Figure 3. PFC neurons differentiated between replays depicting different arms. **A**, Linearized trajectory of a rat running across the three arms of the Y-maze. Arrows depict direction of running corresponding to the arrows in Figure 1B. **B**, Raster of simultaneously recorded HP neurons' firing, ordered by their preferred position on the track. **C**, Estimate of position using Bayesian decoding with HP. **D**, Three example ripple events. Top, Raw LFP. Middle, Ripple filtered LFP. Bottom, Replays decoded of three different arms. **E**, Replay-triggered raster plot and averaged firing rates (**F**) for three example PFC neurons, color coded by the arm identity of the replay. Red and black bars in **F** show baseline and modulation periods. **G**, Replay-triggered firing rates for five more example neurons. **H**, right, A significant number of PFC neurons were significantly differentially

Specifically, the normalized firing rate across each arm (mean firing rate on each arm divided by the sum of mean firing rates across all three arms) was compared the modulation to each arm-replay normalized in the same way [the mean modulation (which is defined above) by each arm-replay was divided by the sum of the mean modulation by each arm-replay]. The absolute value of that difference was summed across all three arms for 4000 permutations of shuffled arm-replay labels and the permutation that minimized the error between the pattern of firing rate across position and pattern of modulation to replay was chosen. This produced a simulated scenario whereby the pattern of modulation a PFC neuron showed to arm-replay matched the pattern of modulation it showed across arms during running. This procedure kept the proportions of arm-replays constant and the way a PFC neuron responded to arm-replay events depicting certain arms and the way it is differentially modulated along arms during running. We then used this permutation of arm-replay labels too test whether this simulated modulation was significantly arm-replay modulated using the same tests that were performed as described above (see PFC arm-replay modulation). To test whether the "Expected" (simulated) modulation could account for the "Observed" selectivity, the sum of squared differences of modulation across arms for the expected and observed replay modulation was compared across neurons (Fig. 4D–F). We also repeated this analysis but using the firing rates across different portions of behavior separately. Each period of time in which the rat fully explored the whole track (each of the three arms in both direction) we computed the average firing rate across the three arms. We then performed the same analysis as described above except using each of these periods of exploration. For each neuron we then took the largest Expected sum of squared differences from these periods of exploration and compared those values to their Observed sum of squared differences and reported these results in the Results.

Arm-replay discriminant analysis using PFC neurons

SWR were classified into arm-replays from the simultaneously recorded PFC neurons with discriminant analysis. For each SWR, the modulation (using the same windows as above, see PFC arm-replay modulation) was calculated for each PFC neuron (excluding those with too few spikes as stated in the previous section). For each session, a classifier on the population modulation (a vector for each PFC neuron's modulation across each arm-replay) was used to discriminate which of the three arms the SWR depicted. In a leave-one-out procedure, the classifier was trained on all but one SWR and tested on that SWR. SWRs were excluded from training or testing whether fewer than four PFC neurons spiked during either the baseline or the modulation window. The MATLAB "classify" function was used, with a *diaglinear* discriminant function, similar to a linear discriminant analysis except with a diagonal covariance matrix estimate. To test significance, the same procedure was done

modulated by replays of different arms (18.5% of PFC neurons were significantly modulated, 28/151; binomial test $p = 2.2 \times 10^{-9}$). Inset, Direction of modulation to ripple events. Left, To control for the possibility of double counting the same neurons in different sessions, we repeated these analyses using the single session from each rat with the highest PFC neuron yield and found similar and significant results (20% of PFC neurons were significantly modulated, 14/69; binomial test, $p = 6.9 \times 10^{-6}$); $***p < 0.001$, $**p < 0.01$, $*p < 0.03$. Shading represents SEM.

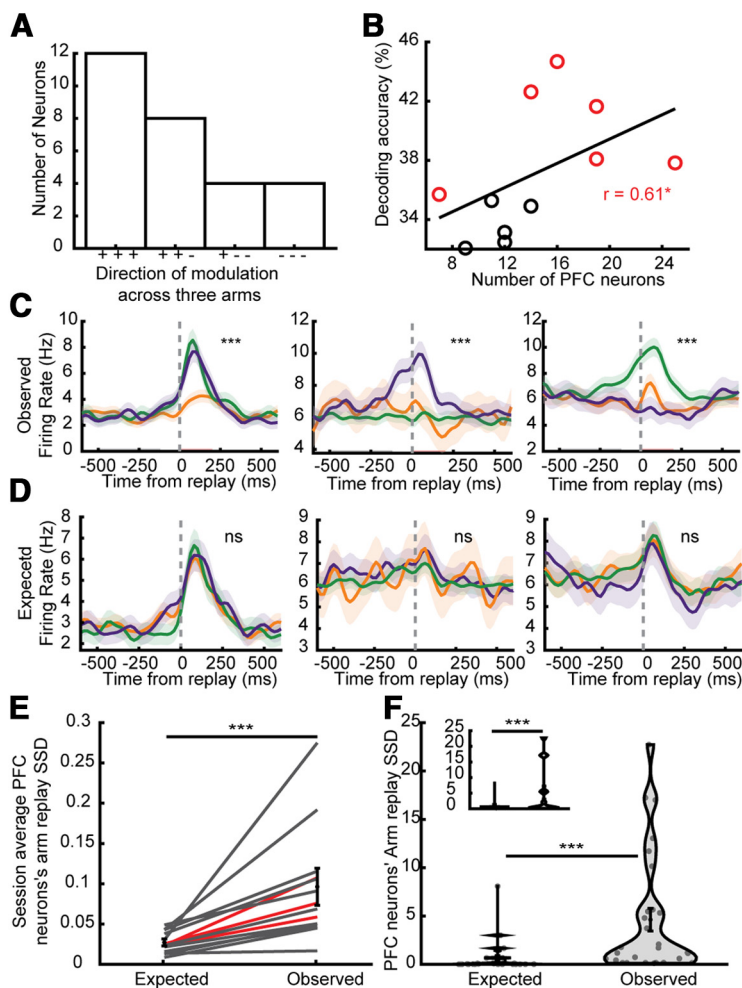


Figure 4. PFC neurons were more modulated to Y-maze arms during replay than behavior. **A**, Number of neurons (of those significantly modulated by arm-replays) that were modulated positively (+) or negatively (-) to replays of each of the three arms. **B**, The accuracy of classifying each ripple into which arm is being depicted in the HP replay using the firing rates of simultaneously recorded PFC neurons was higher with more simultaneously recorded PFC neurons. Spearman's correlation, $N = 11$, $R = 0.62$, $p = 0.045$. Spearman correlation between number of neurons and permutation test p value: $R = -0.64$, $p = 0.033$. Each session that was significantly better than a shuffled distribution (permutation test $p < 0.05$) is plotted in red (all data combined, $p = 0.002$). **C**, Replay-triggered firing rates of the neurons depicted in Figure 2E. **D**, Simulated expected response of the same example PFC neurons as in **A** if they were to respond to arm-replays in a similar ratio of firing as while running (all $p > 0.2$). **E**, Sum of squared differences (armSSD) of modulation across replay of different arms across all PFC neurons for all sessions (black; red lines are averaged for each rat). The armSSD in Observed modulation (examples in **A**) was greater than Expected (examples in **B**). Two-sided Wilcoxon signed-rank test, $N = 11$, 11, signed-rank = 66, $p = 9.8 \times 10^{-4}$. Error bars represent SEM. **F**, The armSSD in Observed were also more than Expected when evaluating only significantly arm-replay modulated neurons. Two-sided Wilcoxon signed-rank test, $N = 28$, 28, $z = 4.44$, $p = 10^{-5}$. Inset, Using only one session per rat. Two-sided Wilcoxon signed-rank test, $N = 14$, 14, $W = 105$, $p = 1.2 \times 10^{-4}$; *** $p < 0.001$, * $p < 0.05$. All error bars and shading represent SEM.

500 times with permuted training labels (which arm was depicted in that arm-replay). The significance of each session and cumulatively across all sessions was assessed by whether the classifier outperformed 95% of the shuffled classifier (permutation test). Additional controls were performed to account for information that (1) ripple power (2) time within the session (3) position of the rat at the time of SWR or (4) the length of the replay, could provide (all $p = 0.002$). To control for ripple power, training and testing was performed separately on the events with the largest and smallest (a median split) mean ripple power. To control for the time during the session, the first half and second half of each session were trained and tested separately using the same procedure. To control of the position of the rat, training and testing was performed separately for SWRs where the rat was on each of the three arms using the same procedure. To control for replay length (in time),

training and testing was performed separately on the events with the longest and shortest (a median split) replays. Although the arm-replay label shuffle keeps the proportions of arm-replays constant and tests for significance given these proportions, we also performed another control to make sure the unequal number of arm-replays was not creating a spurious effect. We repeated the discriminant analysis using a down-sampled set of arm-replays with an equal number for each arm. The results of discrimination using this down-sample control was also significantly better than chance ($p = 0.002$).

Local and non-local internal HP representations during running

All running time (speed > 5 cm/s) was decoded using 40-ms bins with a 10-ms sliding window using HP putative pyramidal neurons and two-dimensional place fields with 8-cm bins smoothed with a Gaussian filter with a radius of 4 cm and a σ of 2 cm (Figs. 1B, C, 6B). To ensure accurate estimates of position, only bins with at least four unique HP neurons fired and five total spikes occurred were included (42.12% of the time bins passed these criteria). If the decoded estimate of position was less than half a position bin away, the bin was considered local, whereas if it was more than half a position bin away it was considered non-local. We also repeated these analyses using an error cut off of 12 and 20 cm, and found similar and significant results. However, a cutoff of 36 cm did not yield similar and significant results. We also repeated these analyses using 5-cm bins (with 5- and 7-cm cut-offs) and 6-cm bins (with 6- and 9-cm cut-offs) and found similar and significant results. PFC firing rates for each decoded bin was taken from the 40 ms after that bin (40-ms delay). We also repeated these analyses with delays of 10–70 ms. PFC firing rate maps as a function of local and non-local HP representation were computed by dividing the summed firing rate from all bins of each position bin by the cumulative time in which that position was decoded (Fig. 6E,F). The same procedure as just described was also performed on time bins when the animal was still (speed < 5 cm/s) to obtain a “replay map” (Fig. 6A, different from the insets which show the modulation to arm-replays as described previously) to use to correlate to two-dimensional local and non-local running maps (Fig. 6G). To see whether local or non-local periods drive PFC to be modulated similarly to the replay map, we computed a Pearson's correlation between either the replay map and the local ($\text{Rho}_{\text{Local}}$) or non-local ($\text{Rho}_{\text{Non-Local}}$) running map and took the difference ($\text{Rho}_{\text{Non-Local}} - \text{Rho}_{\text{Local}}$) across all PFC neurons (Fig. 6H) or only the PFC neurons that were significantly differentially modulated by arm-replays (Fig. 6I).

Basic theta analysis

The LFP from the HP tetrode with the largest number of simultaneously recorded pyramidal neurons was filtered through 6- to 12-Hz bandpass filter, a Hilbert transform was used to obtain the phase, and this phase was adjusted to a global zero phase based on the maximal HP pyramidal neuron activity. PFC neuron locking to HP theta was assessed using any-time the rats speed exceeded 10 cm/s and a Rayleigh permutation test was used to test for significance and the mean resultant vector length was used to assess correlation with ripple modulation (Fig. 2).

For the rest of theta analysis theta was defined with the following procedure. The LFP from the HP tetrode with the largest number of simultaneously recorded pyramidal neurons was filtered through 6- to

12-Hz bandpass filter, a Hilbert transform was used to obtain the phase. As we have done previously (Feng et al., 2015), we then found the peak of theta (360°) and interpolated the spikes to find the phase of theta for each cycle. Theta sequences were then decoded for all sessions and a mean theta sequence was computed and rotated circularly across theta phase to optimize the highest slope of the sequence (Fig. 5A). This shift was applied to all theta cycles in each session for all analysis after Figure 2 to center theta to the theta sweep (Feng et al., 2015). Analyses and results were also similar and significant if each session's optimized theta sequence shift was computed separately instead.

HP-PFC cross-covariance during theta

Standardized cross-covariance during theta was calculated similar to previous reports (Siapas et al., 2005; Jadhav et al., 2016). The cross-covariance between individual PFC neurons and the HP neuron population (only putative pyramidal neurons) was assessed during periods during which the rat was moving (speed >5 cm/s) for at least a second. The cross-covariance was computed separately two ways, using only the HP spikes occurring in the first half of theta cycles or the second half. In order to make sure the same amount of time was being considered for each of these halves, each theta cycle was halved. As we have done previously (Feng et al., 2015), we found the peak of theta (360°) and interpolated the spikes to find the phase of theta for each cycle. This ensures that there is equal amount of time in early as late theta phases. Spikes from theta cycles shorter than 100 ms or longer than 200 ms were not considered. The cross-correlation for each of these conditions was computed for a 500-ms window using 10-ms bins, and each of the conditions was normalized to account for differences in mean firing rates such that the covariances of autocorrelations was one at zero lag (Fig. 5C). The cross-covariance from each of these high velocity periods was normalized by multiplying it by the square-root of the bin size multiplied by the length of the period considered. The cross-covariances for each of these periods was then averaged and smoothed (using a Gaussian of length 50 ms and a SD of 16.7 ms) for each PFC neuron (Fig. 5C). The difference in the cross-covariance between the first and second half of the theta cycle was computed by taking the maximum covariance in the 0- to 250-ms (where HP spikes lead PFC spikes) window of the cross-covariance. The maximum covariance between the first and second half of the theta cycle was compared across neurons with a two-sided Wilcoxon signed-rank test. The difference between the maximum cross-covariance driven by second half spikes subtracted by that of first-half spikes was compared (across PFC neurons) to the extent of arm-modulation (armSSD, described above) using a Pearson's correlation (Fig. 5D). This difference was also compared with the extent of modulation by SWRs (SWR modulation, described previously) with a Pearson's correlation. In addition, we performed a partial correlation between the theta difference and armSSD of PFC cells, controlling for the SWR modulation.

Local and non-local HP portions of theta cycle during running

In these analyses, and all analyses where Bayesian decoding was performed on a theta periods, time bins were used for decoding, based off of theta cycles identified as described earlier (see above, Basic theta analysis). Each theta cycle was assumed to be 120 ms long, starting 120 ms before the end of each theta cycle. Theta cycles during which the rat was running (>5 cm/s) were considered, excluding any time periods

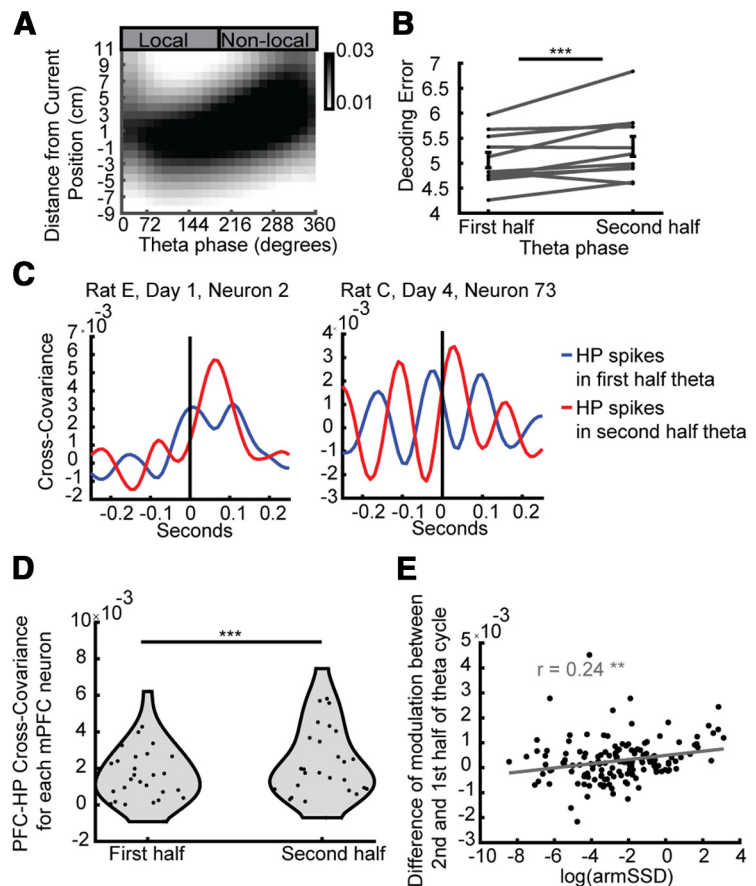


Figure 5. Late-phase HP spikes were more non-local and more closely related to PFC neurons' spiking. **A**, Average theta sequence (decoded estimate of position from HP neurons across HP theta phases, centered on the current position of the rat). **B**, The second half of the theta sequence cycle was more non-local than the first. One-sided Wilcoxon signed-rank test, $N = 11, 11, z = -2.36, p < 0.007$. **C**, Example PFC neurons' (left, rat E, day 1, neuron 2; right, rat C, day 4, neuron 73) cross-covariance with the HP spikes from either the first or second half of theta cycles. Positive values of the x-axis correspond to HP spikes leading PFC spikes. In both neurons the peak of the cross-covariance was higher when using HP spikes from the second half of theta. **D**, Across all significantly differentially modulated PFC neurons, the peak cross-covariance driven by HP spikes was higher for spikes coming from the second half of the theta cycle than the first. Two-sided Wilcoxon signed-rank test, $N = 28, N = 28, p = 6.8e-5$. **E**, Across all PFC neurons, the more differentially modulated a neuron was (log of the sum of squared differences (armSSD) of modulation between replays of different arms) the higher the cross-covariance peak related to HP spikes in later theta phases compared with earlier phases of theta was. There are significant Pearson's correlations between both the arm SSD ($R = 0.32, p = 7.6e-5$) and the log of the armSSD ($R = 0.24, p = 3.6e-3$); *** $p < 0.001$, ** $p < 0.01$. Error bars represent SEM.

overlapping with SWRs. In order to have confidence of the HP estimate of position, only theta cycles with at least five unique HP neurons participating were used. Two portions of each theta cycle were then considered, the local (−120 to −60 ms from the end of the theta cycle) and non-local (−60 to 0 ms from the end of the theta cycle) and Bayesian decoding (using linear fields of HP putative pyramidal neurons) was performed to determine the estimate of position. Portions of theta cycles were not included in the analysis if there were fewer than four spikes or if fewer than two HP neurons spiked. The difference between the estimated and correct estimate of position was calculated and compared between the local and non-local parts of the theta cycle (Fig. 5B).

Similarly to the maps generated above (see Local and non-local internal HP representations during running) the PFC firing rate map was determined for each neuron as a function of decoded position, except here using firing rates from either the 60 ms after the end of each HP theta cycle (non-local; Fig. 6M, right) or the 60 ms before the end of the HP theta cycle (local; Fig. 6M, left), regardless of the decoded content within the theta cycle.

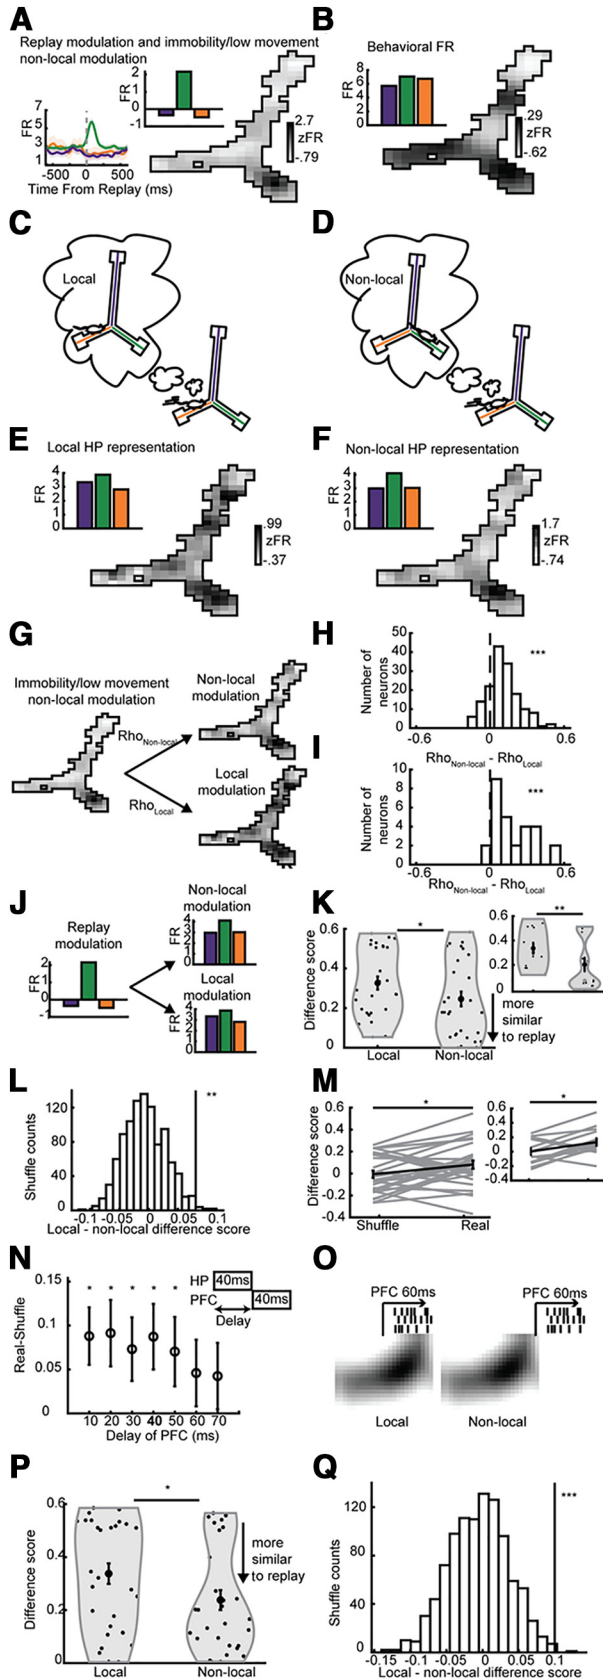


Figure 6. PFC neurons were modulated to non-local HP representations similarly during rest and running. All panels show PFC neurons’ firing patterns as a function of conditions determined by HP firing. **A**, A PFC neuron’s modulation to non-local HP representations of position while the rat is still. Bottom inset, Neuron’s arm-replay triggered firing rate. Top insets in **A**, **B**, **F**, **G** show bars depicting the PFC neurons’ firing pattern across the arms in

Arm-specific difference score

The pattern of modulation to arms depicted during different running conditions [for both the HP internal representation described above (see Local and non-local internal HP representations during running) and the HP theta cycle analysis described above (see Local and non-local HP portions of theta cycle during running)] was compared with the pattern of modulation to arms depicted in replay (Figs. 2, 3, 6A, insets). The following procedure denoting “local” and “non-local” was done using either the local and non-local modulations calculated from the internal representation (Fig. 6E,F,I) or from the local and non-local portions of the HP theta cycle (Fig. 6M). The firing rates as a function of local or non-local position across each arm was averaged (Fig. 6E,F, insets). Then, the replay modulation pattern (Fig. 6A, insets) was compared with the arm-averaged local and non-local modulation using a difference score computed with the following procedure. The modulation values for each condition (replay, and either local or non-local running) normalized $([value - \max(values)] / ([\min(values) - \max(values)])$ across the three arms and all combinations of normalized arm values were compared by taking the absolute value of the difference (e.g., $Arm1_{Replay} - Arm1_{Local}$, $Arm1_{Replay} - Arm2_{Local}$, $Arm1_{Replay} - Arm3_{Local}$, $Arm2_{Replay} - Arm1_{Local}$, etc.). Then the matched (e.g., $Arm1_{Replay} - Arm1_{Local}$) combinations (the diagonal) were summed and divided by the sum of all combinations. This difference score results in smaller values when the pattern of modulation is more similar to replay. Then the difference score for local and non-local was compared across all PFC neurons that were significantly differentially modulated to arm-replays (Fig. 6K,P).

HP internal representation analysis shuffle

To make sure the similarity between PFC modulation to high speed non-local HP representation and arm-replays was not an artifact of a bias in which locations tend to be represented non-locally, we generated a distribution of non-local difference scores using 1000 shuffles of the non-local decoded positions (which shuffles the arm-identity of the positions) and compared the effect (local difference score – non-local difference score; Fig. 6K) to the shuffled distribution of the same calculation using a permutation test (Fig. 6L). We also compared the real effect for each neuron to that neuron’s shuffle distribution with a one sided Wilcoxon signed-rank test (Fig. 6M,N).

HP theta cycle analysis shuffle

To access the significance of the effect of local minus non-local difference score when using portions of the HP theta cycle (Fig. 6P),

each of the four conditions. **B**, The same neuron’s firing rate across position of the rat during running. HP local (**C**) and non-local HP (**D**) representation during running. PFC firing rate map as a function of local (**E**) and non-local (**F**) HP representation. **G**, Depiction of Pearson’s correlation between the PFC replay map (as in **A**) with the local (Rho_{Local}) and non-local ($Rho_{Non-local}$) maps (as in **E**, **F**). **H**, The difference between non-local and local correlations ($Rho_{Non-local} - Rho_{Local}$) across all PFC neurons. One-sided Wilcoxon signed-rank test, $N = 158$, $z = 8.53$, $p < 10^{-17}$. **I**, Across significantly arm-modulated PFC neurons. One-sided Wilcoxon signed-rank test, $N = 28$, $z = 4.41$, $p < 10^{-5}$. **J**, A difference score was computed between PFC modulation to arms during replay and both local and non-local modulation to arms. **K**, The non-local difference score was less than that of local, one-sided Wilcoxon signed-rank test, $N = 28$, $z = 2.06$, $p = 0.0197$. Inset, Only using one session per rat, $N = 14$, $z = 14$, signed-rank = 93, $p = 4.3e-3$. **L**, The true local-non-local difference was greater than a distribution produced by shuffling non-local representations, $p = 9.0e-3$. **M**, The real difference score was greater than the shuffle across the significantly arm-replay modulated neurons, one-sided Wilcoxon signed-rank test, $N = 28$, $z = -2.1$, $p = 0.017$. Inset, Only using one session per rat, $N = 14$, $z = 14$, signed-rank = 18, $p = 0.015$. **N**, The analysis presented here used a 40ms delay between HP local/non-local activity and PFC activity (inset and bolded x-axis label), but delays between 10 and 50 ms also produced similar and significant results; $p = 0.014$, 0.012, 0.021, 0.017, 0.033, 0.12, 0.14. **O**, A schematic showing the comparison of PFC’s response to local versus non-local parts of the theta cycle. **P**, Non-local portions of the theta cycle produced PFC firing patterns across the three arms that better matched PFC’s responses to replay. One-sided Wilcoxon signed-rank test, $N = 28$, $z = 1.88$, $p < 0.04$. **Q**, Permutation test with shuffled labels of first and second half of theta cycle, $p < 0.002$. All error bars represent SEM; *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. Shading represents SEM.

we generated a shuffled distribution of the same calculation with 1000 permutations of shuffled labels of the theta portion (local vs non-local) and compared that shuffled distribution to the true value using a permutation test (Fig. 6Q).

Remote theta cycle analysis

Theta was defined in the same way as above (see Local and non-local HP portions of theta cycle during running), including that only theta cycles during which the rat was running at high speeds (>5 cm/s) were considered, excluding any time periods overlapping with SWRs. An estimate of position was calculated using directional Bayesian decoding every 5 ms with a 20-ms window (Fig. 7A). In order to have confidence of the HP estimate of position only 20-ms bins with at least two unique HP neurons spiking were used as estimates of position, and only theta cycles with at least five unique HP neurons participating were considered. The arm that the animal occupied was decoded using the inbound fields while the other two arms were decoded using outward fields (representing the map of the actual possible trajectory of an inbound rat). For each theta cycle during which the rat was running toward the center of the Y-maze (inbound), the proportion of the posterior in each of the outbound arms was summed for each time bin, and normalized by the length of each arm to produce a representation proportion. The remote arm with the maximum representation was found for each theta cycle.

To see whether these theta cycles were more likely to depict the arm the rat chose next, the proportion of theta cycles depicting the chosen and unchosen arm (no difference = 0.5) was calculated for each lap. As a control, we performed 1000 shuffles where the chosen arm and unchosen arm was permuted within trajectories starting on each arm (so the proportion of choices the rat took on each arm to every other arm was constant, and the proportion of arms that were represented in the remote theta cycles was constant but the relationship between the theta cycles and choice was shuffled). The remote theta cycles were considered to be representing the animal's choice of next arm if the proportion of theta cycles representing that arm exceeded 95% of the shuffled distribution (Fig. 7B, permutation test). For any given starting arm, if the rat did not choose one of the other two arms enough to have an adequate enough distribution of choices (if the number of combinations was less than the number of shuffles) that arm was skipped for that session. Two of the 11 sessions were not used because no arms were viable. Three sessions had three viable arms, two sessions had two viable arms, and four sessions had one viable arm.

Choice discriminant analysis using PFC spikes

To assess whether the PFC predicts the rat's upcoming choice during theta, we performed discriminant analysis using the modulation of PFC neurons to each theta cycle. For each starting arm in each session, a classifier on the population modulation (a vector for each PFC neuron's modulation across each theta cycle) was used to discriminate which of the two arms the rat would choose. Separately for each of the three arms, in a leave-one-out procedure, the classifier was trained on all but one lap and tested on that lap. The MATLAB "classify" function was used, with a "diaglinear" discriminant function. As a control, the same procedure was done 1000 times with permuted training labels (the chosen arm) for each theta cycle. The significance of each session and cumulative significance across all sessions was assessed by whether the classifier outperformed 95%

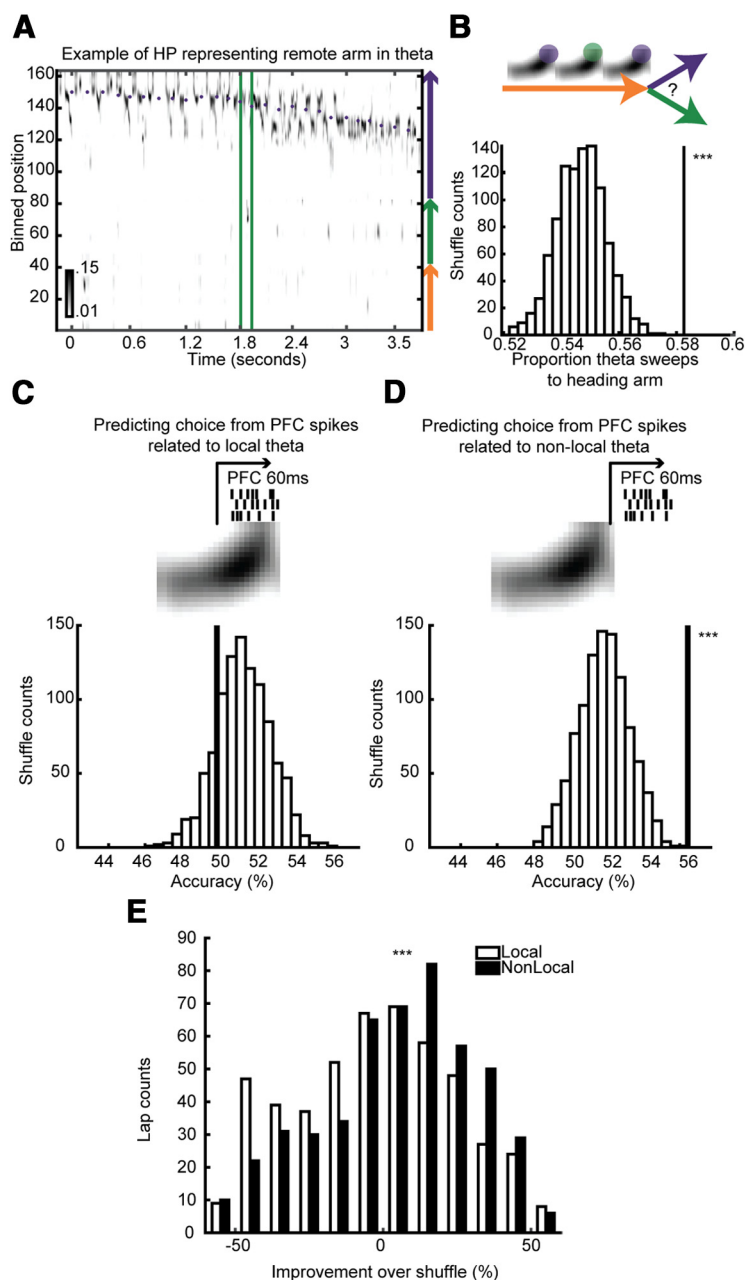


Figure 7. Both HP non-local representations and PFC spiking predicted upcoming choice better than chance during running. **A**, Example of a theta cycle with a representation of a non-local arm. **B**, top, We asked whether the proportion of remote representations during theta tended to overrepresent the arm the rat would choose. Bottom, The proportion of theta sweeps that match the rat's choice, compared with a distribution of shuffled arm labels. Permutation test, $p < 10^{-3}$. **C**, top, We asked whether PFC neurons' simultaneous firing rates could discriminate which arm the rat would choose from spikes directly following the local or (**D**) non-local part of HP theta cycles. Bottom, The accuracy of discriminating the rat's next choice (line) compared with a distribution of accuracies from shuffling the label of the choice. Spikes from non-local (**D**, $p < 10^{-3}$) but not local (**C**, $p > 0.8$) were significantly able to discriminate the rat's choice. **E**, The improvement over the shuffled distribution was compared between non-local and local showing non-local PFC spikes discriminated future choice better than local. One-sided Wilcoxon rank-sum test, $p < 10^{-4}$.

of the shuffled classifier (permutation test). This procedure was done separately using either the 60 ms after the end of each HP theta cycle (non-local; Fig. 7D) or the 60 ms before the end of the HP theta cycle (local; Fig. 7C). To assess whether the non-local classification was more predictive of choice than the local classification (Fig. 7E) we compared the improvement over shuffle for each lap for local and non-local classification with a one-sided Wilcoxon rank-sum test (Fig. 7E). Importantly, this method of comparing the two

conditions does not change in its rigor as a function of how many shuffles are performed to build the control distributions. Arms with fewer than six laps were not used, and laps without an adequate number of different non-local theta representations (if the number of combinations was less than the number of shuffles) were not used.

Results

Spatial responses of prefrontal and HP neurons recorded during learning

HP and PFC neurons (Fig. 1A) were recorded simultaneously from four rats in 11 recording sessions as they explored an asymmetric Y-maze (Fig. 1B). Recordings were taken from learning days 1–5. The rats were rewarded in accordance to an alternation rule and all recordings were performed before asymptotic performance (Fig. 1C). On the 11 recording days, rats had an average reward rate of 63.00% (26.32–78.00%) and average alternation accuracy of 45.57% (0–63.64%). Overall, rats increased their reward rate between the first and last session, from 59.98% to 70.51% (Kendall's rank correlation between recording day and reward rate, $R = 0.51$, $p = 1.2e-3$), and increased their alternation accuracy from 31.32% to 58.32%. These results are consistent with the hypothesis that the rats were learning from their experiences during the time that we were recording their neural activity.

In each of the 11 sessions, we recorded simultaneously from HP neurons (mean = 88, range 62–137, 968 total) and PFC neurons (mean = 14.4, range 7–25, 158 total). Consistent with previous reports, while HP neurons tended to fire in specific regions of the track known as place fields (Fig. 1D), PFC neurons' firing patterns were much less spatially localized (Fig. 1E; Jung et al., 1998; Zielinski et al., 2019). HP neurons exhibited more spatial selectivity (Fig. 1E) and less radial symmetry (Fig. 1F) than PFC neurons. Thus, PFC neurons were less selective for the rat's position during running than HP neurons were. To explore what behavioral correlates better fit PFC neurons' firing rates during running we fit a GLM and using 10-fold cross-validation, predicted seven behavioral variables. We found that across many different time bin sizes, the best predicted variable was the distance traveled within a trial (Fig. 1G). This was better predicted than the time elapsed during the trial, distance from the center of the maze, arm, speed, correctness of the trial, or direction traveled (in/out).

PFC neurons were modulated by HP ripple events and HP theta

In each session, we identified SWRs while the rat was paused on the track (mean = 1985.3, range: 696–2909, total: 21,838; Fig. 2D). Similar to previous reports (Jadhav et al., 2016; Tang et al., 2017; Yu et al., 2017), we found that 57% of PFC neurons (68 neurons) were significantly modulated by SWRs (Fig. 2A,B,D). Of those, 43% were positively modulated, and 57% negatively modulated. When rats run there is a strong theta rhythm (~8 Hz) in the HP, and consistent with prior reports (Jones and Wilson, 2005a,b; Benchenane et al., 2010; Jadhav et al., 2016), we found that 47% of PFC neurons were significantly modulated by HP theta (Fig. 2C–E). Although ripple events and theta are different HP states and occur independently during ongoing behavior, the PFC neurons that were modulated by HP ripple events were more likely to be HP theta modulated (58%) than those that were not (26%). Similarly, PFC neurons that were significantly modulated by HP theta were more likely to be modulated by HP ripple events (80%) than those that were not (50%).

Additionally, the extent to which a PFC neuron was modulated by HP ripple events (absolute value of the average modulation) predicted how modulated it was by HP theta (mean resultant vector length; Fig. 2F). As others have reported (Jadhav et al., 2016), these results depict a HP-PFC network that is co-modulated during both awake rest and behavior.

PFC neurons were selective for arm identity within HP replay

To determine whether PFC neurons were modulated by the spatial information in HP ensembles during SWRs, we used Bayesian decoding (see Materials and Methods) to estimate the posterior probability of position expressed by the HP population during SWRs. We found that SWR-related HP activity depicted trajectories along individual arms, as described previously in a study using a subset of this dataset (Wu and Foster, 2014). In order to increase the power in analyses, unlike in our previous study, we use lower criteria for an arm-replay whereby if the event is decoded to represent one of the arms it is designated an arm-replay (mean: 1243.3, range: 455–2048, total: 13,676; Fig. 3D). We then examined whether PFC neurons were differentially modulated by arm-replay depicting different arms. Overall, a significant portion, 18.5% (28/151; binomial test $p = 2.0e-9$) of PFC neurons were differentially modulated by replay of different arms (Fig. 3E–H). For example, the neuron depicted in the left column of Figure 3E,F (rat E, day 1, neuron 2) increased its firing rate in response to HP replay of the left and right arms, but not to HP replay of the center arm. Alternatively, the neuron depicted in the center column (rat C, day 2, neuron 85) increased its firing rate in response to HP replay of the left arm, showed no change in firing rate to HP replay of the right arm, and decreased slightly its firing rate in response to HP replay of the right arm. Out of those PFC neurons that were significantly modulated by SWRs, 19% (17/90; binomial test $p < 10^{-6}$) of them were also differentially modulated by arm. In order to control for the possibility of double counting the same neurons in different sessions, we repeated these analyses using the single session from each rat with the highest PFC neuron yield and found similar and significant results (20% of PFC neurons were significantly modulated, 14/69; binomial test, $p = 6.9e-6$).

There was a diversity of responses of PFC neurons to replay of different arms, with different PFC neurons exhibiting different combinations of preferences (Figs. 3G, 4A). In another analysis to test the consistency of PFC neuron's selectivity to arm-replay, we hypothesized that it might be possible to discriminate which arm was replayed by HP at any moment, from the PFC neurons' responses to these events alone. Indeed, discriminant analysis using the modulation of all simultaneously recorded PFC neurons (not just those deemed statistically significantly modulated by arm-replay) performed better than chance (discriminating with shuffled training labels; permutation test $p = 0.002$). Importantly, the ability to decode the content of HP arm-replay was not dependent on which arm the animal was on, the time that had passed in the recording session when the event occurred, the degree of ripple power within the event or the length of time each replay took (all permutation tests $p = 0.002$; see Materials and Methods). The yield of PFC neurons varied across sessions, and sessions with more simultaneously recorded PFC neurons discriminated between arms with greater accuracy (Fig. 4B). This suggests that the accuracy of decoding the content of HP arm-replay would increase further with more PFC neurons recorded simultaneously.

PFC neurons were more selective for arm identity during HP replay than during behavior

In order to understand this PFC selectivity further, we considered the hypothesis that these PFC firing rate changes reflected replay of previous PFC activity patterns during spatial behavior. This hypothesis follows naturally by analogy with HP replay, in which replay activity patterns are related to activity patterns that occurred during previous behavioral experience, whether directly (Lee and Wilson, 2002) or indirectly (Foster and Wilson, 2006; Gupta et al., 2010; Pfeiffer and Foster, 2013). Indeed, previous studies of cortical activity have offered support for this hypothesis (Ji and Wilson, 2007; Jadhav et al., 2016; Ólafsdóttir et al., 2016; Yu et al., 2018; Shin et al., 2019). Specifically, we hypothesized that PFC activity during HP replay of an arm would resemble the PFC activity that occurred during behavioral exploration of that same arm. In order to test this hypothesis, we simulated PFC activity during SWRs based on the previous behavioral activity patterns, and compared this to the actual PFC activity patterns that occurred during SWRs (Fig. 4D). In generating these simulated responses, we kept the proportion of arm-replay the same, but non-randomly shuffled the response of each PFC neuron to minimize the difference between the pattern of modulation elicited and the pattern the PFC neurons exhibited across running (Fig. 1D). We used the mean firing rate of each arm because the PFC firing patterns along an arm's length were similar across arms (Fig. 1F; Materials and Methods; mean radial symmetry: 0.27; mean overlap measure: 0.78), and furthermore the similarity across arms was even stronger in the PFC neurons that showed arm-replay specificity (radial symmetry: selective $N=28$, mean: 0.39, range: 0.08–0.90; non-selective $N=123$, mean: 0.25, range: -0.31 –0.93; Wilcoxon rank-sum test: $p=7.2e-3$; overlap measure: selective $N=28$, mean: 0.84, range: 0.58–0.95; non-selective $N=123$, mean: 0.78, range: 0.20–0.96; Wilcoxon rank-sum test: $p=0.012$). Contrary to our hypothesis, these simulated responses were significantly less differentially modulated than the true data (Fig. 4A–D). The sum of squared differences in Expected is less than in Observed across all 11 sessions (Fig. 4E) in addition to only in significantly arm-replay modulated neurons (Fig. 4F). Thus, PFC neurons' selectivity to arms during running did not explain their selectivity to arms during replay. We also repeated the same analyses using the PFC firing rates across the three arms separately for each period of time the rat explored all three arms, and computed the Expected sum of squared difference for each of those exploration periods. When we compared the maximum Expected sum of squared differences that was possible across those periods the simulated responses were significantly less than the true data (two-sided Wilcoxon signed-rank test, $N=28$, $z=2.00$, $p=0.045$; using only one session per rat. Two-sided Wilcoxon signed-rank test, $N=14$, $W=90$, $p=0.017$). Thus, not even in certain subsets of behavior did PFC neurons' selectivity to arms during running explain their selectivity to arms during replay.

PFC neurons that were modulated by HP arm-replay were also more active after HP spikes from non-local theta phases

Given the surprising degree of spatial coding exhibited by PFC neurons during HP replay, despite the lack of spatial specificity in the same neurons' responses during running behavior, we wondered whether alternative measures of co-activity between PFC and HP neurons during running would reveal a different basis for coordination between the two areas. We first noted that when rats run through space, the HP internal representation of space could be local or non-local, that is, either reflecting the

current position as expected from classical place field responses, or reflecting the recall of distant place information in a manner analogous to replay, respectively. Non-local and local HP representations corresponded to different parts of "theta sequences" or "theta sweeps," which are sequential representations of space nested within cycles of the strong theta rhythm in the HP during running (Fig. 5A; Foster and Wilson, 2007; Johnson and Redish, 2007). The second half of the theta sequence cycle tended to be more non-local than the first (Fig. 5B). Because PFC neurons were modulated by HP activity during both rest and running (Fig. 2E) as well as being modulated by non-local information in replay (Fig. 3), we next asked whether PFC neurons were active more after non-local parts of HP theta. We measured the cross-covariance of HP spikes and each PFC neuron using HP spikes from either the first or second part of the theta cycle (Fig. 5C), and determined that PFC neurons were more strongly related to HP spikes from the second half of theta cycles both across the whole population of PFC neurons (two-sided Wilcoxon signed-rank test, $N=158$, $N=158$, $p=4.9e-5$), as well as only the PFC neurons that were significantly arm-replay modulated (Fig. 5D). The peak of HP to PFC modulation is a 10-ms lag when using all spikes and 30 ms when using only second half of theta cycle spikes, while the peak using only the first half of the theta cycle has a negative lag of 30 ms. Furthermore, we found that the extent to which a PFC neuron was differentially modulated by replay of different arms predicted how much more modulated the PFC neuron was to HP spikes from the second half of theta cycles compared with the first half (Fig. 5E). There was no significant relationship between the extent a PFC neuron was modulated by SWRs (as in Fig. 2G) and how much more modulated the neuron was to HP spikes from the second half of theta [Pearson's correlation: $R=0.040$, $p=0.62$, $\log(\text{SWR modulation})$: $R=-0.033$, $p=0.68$]. And when taking into account the extent a PFC neuron was modulated by SWRs, there was still a significant correlation between the extent it was driven by late theta spikes and how differentially modulated it was to arm-replays (partial correlation: $R=0.31$, $p=9.9e-5$; log: $R=0.24$, $p=3.5e-3$). Thus, PFC neurons that were modulated by non-local information during rest were also active more after late phases of HP theta, which were often associated with non-local representations during behavior.

PFC neurons' arm preference in replay is best explained by HP non-local representations during run

Since we found that PFC neurons were tuned to the internal HP representation of arm during replay and active more after late phases of theta, we hypothesized that the internal HP estimate of position during non-local representation might be a better predictor of how the PFC neurons respond to replay than the true external position. We decoded the position using 40-ms bins of HP activity every 10 ms of each rat as he ran and measured how PFC neurons' firing rates were modulated 40 ms later by the 2-d spatial representation in HP, broken down by local (Fig. 6C,E) or non-local (Fig. 6D,F) positions. We compared the PFC neurons' non-local modulation map during immobility (Fig. 6A, mirroring the selectivity to arm-replay shown in the insets) to either the local or non-local modulation maps during running for each PFC neuron (Fig. 6G). We found that PFC neurons' immobility non-local modulation was more correlated with the non-local modulation maps than the local modulation maps during running (Fig. 6H). Further, when we restricted these analyses to only neurons that were previously found to be significantly differentially modulated by replay of different arms, we found that this

population's replay modulation showed a strong correlation to non-local compared with local HP activity (Fig. 6J). Because these analyses were done on the 2D map of modulation, they do not necessarily reflect the pattern across the three arms but could also reflect similar modulation patterns within those arms. To specifically test whether the pattern of non-local modulation by different arms that PFC neurons showed to the HP representation during running better reflected the arm-replay modulation than local HP representation, we computed a difference score for each neuron for non-local and local separately, and found that the pattern of modulation across the three arms during replay was better reflected by non-local than local HP representation (Materials and Methods; Fig. 6J,K and 6A,B,E,F, insets). To control for the possibility that some bias in the non-local HP representations happened to produce this effect, we performed a permutation test where the non-local representations were shuffled, and found that the pattern of modulation produced by the true non-local representations was unlikely to occur by chance (Fig. 6L). Across arm-replay modulated neurons the pattern of modulation matched better than chance (Fig. 6M). In addition to using a 40-ms delay (which allows for HP activity from 0–80 ms previous to be reflected in PFC firing, as depicted in Fig. 6N, inset), we tested other delays and, consistent with the results from cross-covariance analysis, found that delays of 10–50 ms produced similar and significant results (Fig. 6N). Thus, during running, PFC neurons exhibited selectivities to non-local HP arm representations that were similar to their individual selectivities to HP replay arm representations, and dissimilar to the HP representation of current position. Hence, whether during running or immobility, PFC neurons were sensitive to the non-local information represented in HP, and insensitive to the actual current location of the animal.

PFC neurons responded similarly to HP replay as they did to remote HP theta cycles

Because PFC neurons' replay selectivity was correlated with their responses to non-local HP representations while running, we asked whether this was true on a theta-cycle-by-theta-cycle level. It is unclear whether theta phase or decoded position is the better way to define non-local activity, so we aimed to use both. Again, we generated modulation maps of PFC neurons by HP representation, but this time instead of comparing non-local and local periods, we compared the first and second half of each theta cycle. Knowing that PFC neurons that were significantly differentially modulated by replay of different arms were active more after non-local parts of theta (Fig. 5D,E), we asked whether these neurons had a pattern of modulation by the non-local part of theta cycles that better matched their replay modulation than the local part of theta cycles (Fig. 6O). We found that the pattern of modulation by the HP representation after the second half of theta (non-local portion) was more similar to their replay modulation pattern than after the first (local) portion (Fig. 6P) and that this effect was greater than chance (Fig. 6Q). Thus, non-local representations in the HP during theta were followed by arm-modulated firing of PFC neurons in a similar pattern as during replay.

Remote theta cycles and associated PFC spiking both predicted rats' choices

Non-local representations in theta sometimes predict the animal's upcoming behavior (Wikenheiser and Redish, 2015) and sometimes do not (Kay et al., 2020). We asked whether, as the rat approached a choice point, the proportion of theta cycles in

which the HP remotely represented each of the two potential next arms (Fig. 7A) predicted which arm it would choose. We found that there was a bias of remote theta cycles toward the chosen arm (Fig. 7B). This bias was not because of there being matched biases of which arm the rat tended to choose and which arm theta cycles tended to represent. This was tested by comparing the bias to a control distribution where the chosen and unchosen arm were shuffled within trajectories starting on each arm, so that the proportion of choices the rat took on each arm to every other arm was constant, and the proportion of arms that tended to be represented in the remote theta cycles was constant, but the relationship between the theta cycles and choice was shuffled (permutation test, $p < 10^{-3}$). Thus, remote theta cycles as the rat ran down one arm tended to reflect the arm the rat was about to choose next.

Because PFC neurons activity was more related to these remote theta cycles (Fig. 6) as well as to HP spikes in later, non-local phases of HP theta (Fig. 5), we hypothesized that PFC spikes might predict the rat's choice. We found that when we attempted to classify which arm the rat would choose on each lap using all PFC neurons' spiking triggered off of the local part of HP theta (Fig. 7C, top) that we could predict the rats' behavior no better than chance (Fig. 7C, bottom). However, when we used PFC neurons' spiking triggered off of the non-local part of HP theta (Fig. 7D, top), we were able to decode the rats' future choice better than chance (Fig. 7D, bottom). Furthermore, the improvement over shuffle using non-locally triggered PFC spikes was significantly greater than the improvement over shuffle using locally triggered PFC spikes (Fig. 7E). Thus, even early in learning, PFC neurons expressed information about the animal's choice, and that information was yoked to the non-local representations in HP theta.

Discussion

We have shown that a statistically significant portion of prefrontal (PFC) neurons can be selective to the spatial trajectory information contained in HP replay. PFC neurons exhibited diverse selectivities for replays of different arms. Because of this diversity and robustness in responses, arm identity of HP replay could be discriminated better than chance from small numbers of simultaneously recorded PFC neurons. Intriguingly, PFC neurons' arm selectivities could not be predicted from those neurons' place responses during behavior. The observed modulation by replayed arm was greater than expected if the PFC neurons had solely been reactivating their firing pattern seen when running across the arms. We found that HP's representations of nonlocal position during running explained the PFC neurons' firing patterns to HP arm-replay better than representations of local position. This suggests that PFC neurons' activities, which were unrelated to the current location of the animal, were consistently tuned to internally-generated non-local representations in the HP.

Previous studies have shown that many areas are modulated by HP ripples (Pennartz et al., 2004; Ji and Wilson, 2007; Peyrache et al., 2009; Logothetis et al., 2012; Jadhav et al., 2016; Ólafsdóttir et al., 2016; Sosa et al., 2020) and that PFC neurons' response to SWRs are modulated by state or context (Tang et al., 2017; Yu et al., 2017). Here, we have shown, for the first time to our knowledge, that representations of specific arms in a maze are reflected in individual cortical neurons' firing rates during SWRs. This finding provides strong support for the hypothesis that the role of HP SWRs is not solely to compute locally within




Figure 8. A schematic showing that while HP neurons reactivate previous encoded experience, PFC neurons instead consistently reflect non-local representations. Top row, As a rat moves along a trajectory from the center arm up into the right arm, HP neurons with place fields along the trajectory are activated locally (top left panel, open ellipses). Additionally, however, HP neurons with fields on the left arm can be activated non-locally (top left panel, solid ellipse), and it is these non-local activations that PFC neurons are tuned to (top right panel, solid ellipse). Bottom row, During rest periods, HP neurons can exhibit replay of the same trajectory as in the top panel (bottom left panel, solid ellipses). However, this replay will typically recruit different PFC neurons to those that were recruited by the non-local HP activity that occurred during the associated behavior (bottom right panel, solid ellipse, compared with top right panel, solid ellipse). Thus, when the HP is reactivating its representations from a previous behavioral episode, the PFC is typically not reactivating its representations from the same episode.

the HP or to upregulate activity non-specifically in cortical areas, but rather to communicate the information within replay to downstream areas (Buzsáki, 1996; Sutherland and McNaughton, 2000; Teyler and Rudy, 2007).

The communication of information from HP to cortex is a necessary component of models of systems consolidation in which the replay of different memories enables the establishment of corresponding cortical representations (McClelland et al., 1995; Nadel et al., 2000; Sutherland and McNaughton, 2000). However, our data do not align well with the Indexing Theory of consolidation (Teyler and Rudy, 2007), which posits that the HP possesses a pointer to a distributed cortical representation that is established at the time of encoding during behavior, and which is subsequently reactivated repeatedly to promote the slow consolidation of synaptic connections between constituent neurons of the cortical representation. In this framework, the cortical and HP patterns refer to the same external information (i.e., the experience to be remembered), established through quick and slow learning, respectively. In contrast, our data demonstrated a mismatch between HP-PFC coactivity patterns during experience and reactivation, when examined on the basis of the animal's actual position. Instead, the PFC was consistently engaged with non-local representations in the HP. Thus, we propose that PFC neuron activity during HP replay does not in general exhibit reactivation, even while HP neurons do (Fig. 8). This schematic figure demonstrates that, to the degree that PFC neurons respond to non-local activity during behavior, they cannot then be reactivating their own activity from this behavior whenever the HP replays it.

Others have reported HP-PFC co-reactivation of encoded experience, but have not reported that PFC neurons' selectivity for portions of the environment during replay is more so than their selectivity during encoding, as we report here (Jadhav et al.,

2016; Tang et al., 2017). These differences may be because of the large numbers of HP neurons we recorded simultaneously and therefore the large amount of replay we were able to resolve, as well as the fact that our experiments took place in the very early learning stages of a task that was being learned slowly and would never become well learned. Other sources of variability across studies may include the sub-areas or layers of PFC recorded from. Experiments using rigid probes instead of independently moveable tetrodes might address these potential differences. Interestingly, a lack of cortical reactivation during HP replay events is compatible with cortical reactivation that is unrelated to HP replay. Replay across PFC neurons has recently been reported (Kafer et al., 2020). The same group has also reported replay across neurons in entorhinal cortex unrelated to HP replay (O'Neill et al., 2017; although see Ólafsdóttir et al., 2016). Integrating these and our findings leads to the suggestion that both cortical replay, and cortical activity driven by HP replay, may occur, but without the relationship between the two assumed by some consolidation theories.

Prior work has shown that activity in cortical areas before HP SWRs can also be related to HP activity (Rothschild et al., 2017; Chung et al., 2019). In this study we aimed to explore how non-local representations in HP such as replay are reflected in PFC neurons, and because of the evidence in our data that PFC neurons were responding to HP activity, we chose to focus our analysis on PFC's activity during and slightly after HP activity. However, our data does not preclude the notion that information may be traveling in both directions and/or from another shared region. Our data also leaves open the exact computations that PFC neurons perform on incoming HP internal representations.

Theta sequences in the HP have also been shown to be modulated by reward and task structure (Gupta et al., 2012; Wikenheiser and Redish, 2015). HP theta can coordinate activity between HP and PFC, as shown here and in previous studies (Siapas et al., 2005; Benchenane et al., 2010). Furthermore, our findings that PFC neurons are modulated by the information within replay in the same pattern as they respond to non-local portions of theta cycles is consistent with the notion that non-local parts of theta sequences are involved in similar computations as non-local replay (Pezzulo et al., 2017). It is also consistent with the idea that theta cycles have encoding and retrieval portions (Siegle and Wilson, 2014) and extends these ideas to posit that non-local parts of theta sequences engage a similar cortical network as do non-local replay.

Our findings are not in opposition to the general framework that PFC neurons are important for learning abstract rules of a task space (Wallis et al., 2001; Rich and Shapiro, 2007; Wilson et al., 2014). The rats recorded from in this study were early in learning, and it is likely that the rats were still trying to determine the relevant dimensions of the task space to choose an effective strategy. However, whatever coding scheme might later develop in the PFC population to support high performance (Rich and Shapiro, 2009; Powell and Redish, 2016) might be secondary to the early integration of HP representations of space. Specifically, having information about the possible remote trajectories that defined the space might be useful to determine patterns or rules. Possibly, with increased experience in this task, the PFC neurons could become more tuned to another dimension of the task and less tuned to HP internal place representation, or begin to generalize across the environment (Tang et al., 2017; Yu et al., 2018).

PFC neurons in our recordings were tuned to similar internal HP representations regardless of the state the rat was in. This is consistent with the notion that PFC is oriented inward, driven

more by internal states than by external stimuli (Brincat and Miller, 2016) and part of a “default mode network” (Schacter et al., 2007; Kaplan et al., 2016). Our data support the idea that PFC performs computations on internal variables: the relationship between PFC neurons’ modulation by replay and firing on the track emerged only given the internal variable of remote positions represented by HP.

Consistent with previous findings of HP and PFC’s involvement in prospective coding (Ito et al., 2015), we found that portions of remote-representing theta cycles, as well as associated PFC spikes, predicted the choice of the rat. Thus, even in early learning when robust “splitter cells” had not yet developed, HP and PFC reflected upcoming choices on the same theta-cycle timescale. PFC receives input from many other brain areas in addition to the HP including the striatum and amygdala (Sotres-Bayon and Quirk, 2010; Catanese et al., 2016), and might use theta to gate its input from different areas (Hyman et al., 2011). This theta gating might provide a mechanism for PFC to tune in selectively to nonlocal representations in HP. In this framework, the PFC would be tuning into internal HP representations during late HP theta phases, especially in periods when HP theta phase is a strong driver of spike timing. Then, in other phases or time periods it might tune into other internal variables that are also relevant for behavior. Further study of how and when PFC neurons are selective to different specific internal variables in the brain will likely provide further insight into PFC function, as well as the mechanisms of remote representations underlying memory and prediction.

References

- Battaglia FP, Sutherland GR, McNaughton BL (2004) Local sensory cues and place cell directionality: additional evidence of prospective coding in the hippocampus. *J Neurosci* 24:4541–4550.
- Benchenane K, Peyrache A, Khamassi M, Tierney PL, Gioanni Y, Battaglia FP, Wiener SI (2010) Coherent theta oscillations and reorganization of spike timing in the hippocampal–prefrontal network upon learning. *Neuron* 66:921–936.
- Brincat SL, Miller EK (2016) Prefrontal cortex networks shift from external to internal modes during learning. *J Neurosci* 36:9739–9754.
- Buzsáki G (1996) The hippocampo–neocortical dialogue. *Cereb Cortex* 6:81–92.
- Catanese J, Carmichael JE, van der Meer MAA (2016) Low- and high-gamma oscillations deviate in opposite directions from zero-phase synchrony in the limbic corticostriatal loop. *J Neurophysiol* 116:5–17.
- Chung JE, Joo HR, Fan JL, Liu DF, Barnett AH, Chen S, Geaghan-Breiner C, Karlsson MP, Karlsson M, Lee KY, Liang H, Magland JF, Pebbles JA, Tooker AC, Greengard LF, Tolosa VM, Frank LM (2019) High-density, long-lasting, and multi-region electrophysiological recordings using polymer electrode arrays. *Neuron* 101:21–31.e5.
- Davidson TJ, Kloosterman F, Wilson MA (2009) Hippocampal replay of extended experience. *Neuron* 63:497–507.
- Ego-Stengel V, Wilson MA (2010) Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus* 20:1–10.
- Feng T, Silva D, Foster DJ (2015) Dissociation between the experience-dependent development of hippocampal theta sequences and single-trial phase precession. *J Neurosci* 35:4890–4902.
- Foster DJ (2017) Replay comes of age. *Annu Rev Neurosci* 40:581–602.
- Foster DJ, Wilson MA (2006) Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* 440:680–683.
- Foster D, Wilson M (2007) Hippocampal theta sequences. *Hippocampus* 17:1093–1099.
- Gupta AS, van der Meer MAA, Touretzky DS, Redish AD (2010) Hippocampal replay is not a simple function of experience. *Neuron* 65:695–705.
- Gupta AS, van der Meer MAA, Touretzky DS, Redish AD (2012) Segmentation of spatial experience by hippocampal θ sequences. *Nat Neurosci* 15:1032–1039.
- Hyman JM, Hasselmo ME, Seamans JK (2011) What is the functional relevance of prefrontal cortex entrainment to hippocampal theta rhythms? *Front Neurosci* 5:24.
- Ito HT, Zhang SJ, Witter MP, Moser EI, Moser MB (2015) A prefrontal–thalamo–hippocampal circuit for goal-directed spatial navigation. *Nature* 522:50–55.
- Jadhav SP, Kemere C, German PW, Frank LM (2012) Awake hippocampal sharp-wave ripples support spatial memory. *Science* 336:1454–1458.
- Jadhav SP, Rothschild G, Roumis DK, Frank LM (2016) Coordinated excitation and inhibition of prefrontal ensembles during awake hippocampal sharp-wave ripple events. *Neuron* 90:113–127.
- Ji D, Wilson MA (2007) Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat Neurosci* 10:100–107.
- Johnson A, Redish AD (2007) Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci* 27:12176–12189.
- Jones MW, Wilson MA (2005a) Phase precession of medial prefrontal cortical activity relative to the hippocampal theta rhythm. *Hippocampus* 15:867–873.
- Jones MW, Wilson MA (2005b) Theta rhythms coordinate hippocampal–prefrontal interactions in a spatial memory task. *PLoS Biol* 3:e402.
- Jung MW, Qin Y, McNaughton BL, Barnes CA (1998) Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cereb Cortex* 8:437–450.
- Kaefer K, Nardin M, Blahna K, Csicsvari J (2020) Replay of behavioral sequences in the medial prefrontal cortex during rule switching. *Neuron* 106:154–165.e6.
- Kaplan R, Adhikari MH, Hindriks R, Mantini D, Murayama Y, Logothetis NK, Deco G (2016) Hippocampal sharp-wave ripples influence selective activation of the default mode network. *Curr Biol* 26:686–691.
- Karlsson MP, Frank LM (2009) Awake replay of remote experiences in the hippocampus. *Nat Neurosci* 12:913–918.
- Kay K, Chung JE, Sosa M, Schor JS, Karlsson MP, Larkin MC, Liu DF, Frank LM (2020) Constant sub-second cycling between representations of possible futures in the hippocampus. *Cell* 180:552–567.e25.
- Lee AK, Wilson MA (2002) Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* 36:1183–1194.
- Logothetis NK, Eschenko O, Murayama Y, Augath M, Steudel T, Evrard HC, Besserve M, Oeltermann A (2012) Hippocampal–cortical interaction during periods of subcortical silence. *Nature* 491:547–553.
- McClelland JL, McNaughton BL, O’Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev* 102:419–457.
- Miller E, Cohen J (2001) An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167–202.
- Nádasy Z, Hirase H, Czurkó A, Csicsvari J, Buzsáki G (1999) Replay and time compression of recurring spike sequences in the hippocampus. *J Neurosci* 19:9497–9507.
- Nadel L, Samsonovich A, Ryan L, Moscovitch M (2000) Multiple trace theory of human memory: computational, neuroimaging, and neuropsychological results. *Hippocampus* 10:352–368.
- O’Keefe J (1976) Place units in the hippocampus of the freely moving rat. *Exp Neurol* 51:78–109.
- O’Neill J, Boccarda CN, Stella F, Schoenenberger P, Csicsvari J (2017) Superficial layers of the medial entorhinal cortex replay independently of the hippocampus. *Science* 355:184–188.
- Ólafsdóttir HF, Carpenter F, Barry C (2016) Coordinated grid and place cell replay during rest. *Nat Neurosci* 19:792–794.
- Pennartz CMA, Lee E, Verheul J, Lipa P, Barnes CA, McNaughton BL (2004) The ventral striatum in off-line processing: ensemble reactivation during sleep and modulation by hippocampal ripples. *J Neurosci* 24:6446–6456.
- Peyrache A, Khamassi M, Benchenane K, Wiener SI, Battaglia FP (2009) Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nat Neurosci* 12:919–926.
- Pezzulo G, Kemere C, van der Meer MAA (2017) Internally generated hippocampal sequences as a vantage point to probe future-oriented cognition. *Ann NY Acad Sci* 1396:144–165.

- Pfeiffer BE, Foster DJ (2013) Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* 497:74–79.
- Powell NJ, Redish AD (2016) Representational changes of latent strategies in rat medial prefrontal cortex precede changes in behaviour. *Nat Commun* 7:12830.
- Rich EL, Shapiro ML (2007) Prelimbic/infralimbic inactivation impairs memory for multiple task switches, but not flexible selection of familiar tasks. *J Neurosci* 27:4747–4755.
- Rich EL, Shapiro M (2009) Rat prefrontal cortical neurons selectively code strategy switches. *J Neurosci* 29:7208–7219.
- Rothschild G, Eban E, Frank LM (2017) A cortical–hippocampal–cortical loop of information processing during memory consolidation. *Nat Neurosci* 20:251–259.
- Schacter DL, Addis DR, Buckner RL (2007) Remembering the past to imagine the future: the prospective brain. *Nat Rev Neurosci* 8:657–661.
- Shin JD, Tang W, Jadhav SP (2019) Dynamics of awake hippocampal–prefrontal replay for spatial learning and memory-guided decision making. *Neuron* 104:1110–1125.e7.
- Siapas AG, Lubenov EV, Wilson MA (2005) Prefrontal phase locking to hippocampal theta oscillations. *Neuron* 46:141–151.
- Siegle JH, Wilson MA (2014) Enhancement of encoding and retrieval functions through theta phase-specific manipulation of hippocampus. *Elife* 3:e03061.
- Singer AC, Carr MF, Karlsson MP, Frank LM (2013) Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. *Neuron* 77:1163–1173.
- Sosa M, Joo HR, Frank LM (2020) Dorsal and ventral hippocampal sharp-wave ripples activate distinct nucleus accumbens networks. *Neuron* 105:725–741.e8.
- Sotres-Bayon F, Quirk GJ (2010) Prefrontal control of fear: more than just extinction. *Curr Opin Neurobiol* 20:231–235.
- Sutherland G, McNaughton BL (2000) Memory trace reactivation in hippocampal and neocortical neuronal ensembles. *Curr Opin Neurobiol* 10:180–186.
- Tang W, Shin JD, Frank LM, Jadhav SP (2017) Hippocampal–prefrontal reactivation during learning is stronger in awake compared with sleep states. *J Neurosci* 37:11789–11805.
- Teyler T, Rudy J (2007) The hippocampal indexing theory and episodic memory: updating the index. *Hippocampus* 17:1158–1153.
- Vertes RP (2006) Interactions among the medial prefrontal cortex, hippocampus and midline thalamus in emotional and cognitive processing in the rat. *Neuroscience* 142:1–20.
- Wallis JD, Anderson KC, Miller EK (2001) Single neurons in prefrontal cortex encode abstract rules. *Nature* 411:953–956.
- Wikenheiser AM, Redish AD (2015) Hippocampal theta sequences reflect current goals. *Nat Neurosci* 18:289–294.
- Wilson RC, Takahashi YK, Schoenbaum G, Niv Y (2014) Orbitofrontal cortex as a cognitive map of task space. *Neuron* 81:267–279.
- Wu X, Foster DJ (2014) Hippocampal replay captures the unique topological structure of a novel environment. *J Neurosci* 34:6459–6469.
- Yu JY, Kay K, Liu DF, Grossrubatscher I, Loback A, Sosa M, Chung JE, Karlsson MP, Larkin MC, Frank LM (2017) Distinct hippocampal–cortical memory representations for experiences associated with movement versus immobility. *Elife* 6:e27621.
- Yu JY, Liu DF, Loback A, Grossrubatscher I, Frank LM (2018) Specific hippocampal representations are linked to generalized cortical representations in memory. *Nat Commun* 9:2209.
- Zielinski MC, Shin JD, Jadhav SP (2019) Coherent coding of spatial position mediated by theta oscillations in the hippocampus and prefrontal cortex. *J Neurosci* 39:4550–4565.