Neuron

Article

Sleep stages antagonistically modulate reactivation drift

Highlights

- Reactivated CA1 assemblies reorganize during prolonged sleep/rest periods in rats
- Assemblies transition from pre-sleep representations to those seen post sleep
- Slow-wave sleep accelerates assembly drift; REM sleep counteracts it
- Rate-changing pyramidal cells drive the reorganization of reactivated assemblies

Authors

Lars Bollmann, Peter Baracskay,
Federico Stella, Jozsef Csicsvari

Correspondence

federico.stella@donders.ru.nl (F.S.),
jozsef.csicsvari@ist.ac.at (J.C.)

In brief

Bollmann et al. track reactivated CA1 assemblies representing spatial memories during 16–20 h of sleep/rest. Assemblies initially reflect recently learned spatial memories but are gradually transformed into those seen during the memory recall session following rest. Whereas slow-wave sleep accelerates the assembly drift, REM epochs counteract it.

Bollmann et al., 2025, Neuron *113*, 1446–1459
May 7, 2025
https://doi.org/10.1016/j.neuron.2025.02.025

CellPress

CellPress
OPEN ACCESS

Neuron

Article

Sleep stages antagonistically modulate reactivation drift

Lars Bollmann,[1,3] Peter Baracskay,[1,3] Federico Stella,[2,*] and Jozsef Csicsvari[1,4,*]
[1]Institute of Science and Technology (ISTA) Austria, 3400 Klosterneuburg, Austria
[2]Donders Institute for Brain, Cognition and Behaviour, Radboud University, 6525 AJ Nijmegen, the Netherlands
[3]These authors contributed equally
[4]Lead contact
*Correspondence: federico.stella@donders.ru.nl (F.S.), jozsef.csicsvari@ist.ac.at (J.C.)
https://doi.org/10.1016/j.neuron.2025.02.025

SUMMARY

Hippocampal reactivation of waking neuronal assemblies in sleep is a key initial step of systems consolidation. Nevertheless, it is unclear whether reactivated assemblies are static or whether they reorganize gradually over prolonged sleep. We tracked reactivated CA1 assembly patterns over ∼20 h of sleep/rest periods and related them to assemblies seen before or after in a spatial learning paradigm using rats. We found that reactivated assembly patterns were gradually transformed and started to resemble those seen in the subsequent recall session. Periods of rapid eye movement (REM) sleep and non-REM (NREM) had antagonistic roles: whereas NREM accelerated the assembly drift, REM countered it. Moreover, only a subset of rate-changing pyramidal cells contributed to the drift, whereas stable-firing-rate cells maintained unaltered reactivation patterns. Our data suggest that prolonged sleep promotes the spontaneous reorganization of spatial assemblies, which can contribute to daily cognitive map changes or encoding new learning situations.

INTRODUCTION

It has been long established that sleep promotes the recall of previously acquired memories, and there is ample evidence linking sleep to systems consolidation.[1] Furthermore, in many instances, the hippocampus plays a critical role in sleep-associated memory stabilization.[2] A long-standing idea suggests that the hippocampus acts as an intermediate storage area for recently acquired memories, and such memories are reactivated in sleep.[3–5] Subsequently, the reactivation of these memory traces would promote their transfer to other cortical areas for consolidation and long-term storage.[6,7] This hypothesis has received consistent experimental support from a multitude of human and animal studies.[8] Among these, rodent electrophysiological experiments could directly examine the content of neuronal activity during reactivation and test their relationship to memory recall. These showed that the rate of reactivation is increased after the experience of novel learning situations, and the content of reactivation predicts subsequent memory performance.[9–12] Furthermore, suppressing or prolonging sharp-wave ripples (SWRs), a local field potential (LFP) event that usually accompanies hippocampal reactivation, disrupts or facilitates, respectively, subsequent memory recall.[10,13] Finally, content-specific disruption of reactivation events leads to selective impairments in the recall of the disrupted memory.[14]

So far, hippocampal reactivation of neuronal activity has been primarily studied during brief (<1 h) periods of sleep or during waking behavior, with the latter linked to different functions such as decision making or short-term memory. These works may have only examined the initial stages of a more complex sleep reactivation process, which may continue for prolonged periods.[15] Reactivation events can be detected over longer (4–6 h) periods of time[16,17] and, in these studies, neuronal patterns of recent experiences were preferentially reactivated during the entire sleep duration. However, it is not clear whether the neural signature of specific experiences remains unchanged over extended periods or whether it eventually undergoes some form of modification and rearrangement due to early consolidation or other network processes.

Here, we quantified reactivation dynamics during up to 20 h of quiet rest periods and sleep following spatial learning and tested to what degree reactivated patterns represented previous spatial learning patterns or resembled those seen at subsequent memory retrieval. We provide evidence for the transformation of reactivated memory traces and show that many of the transformed reactivated neuronal patterns will recur in the subsequent spatial memory recall trials. We also show that rapid eye movement (REM) and non-REM (NREM) periods exert opposite actions on this transformation. Finally, we show that this transformation is driven by a subgroup of pyramidal cells only and that interneuron activity also reflects the speed of transformation.

RESULTS

We performed 128-channel wireless recordings in the dorsal CA1 region in three Long-Evans rats using bilaterally implanted

1446 Neuron *113*, 1446–1459, May 7, 2025

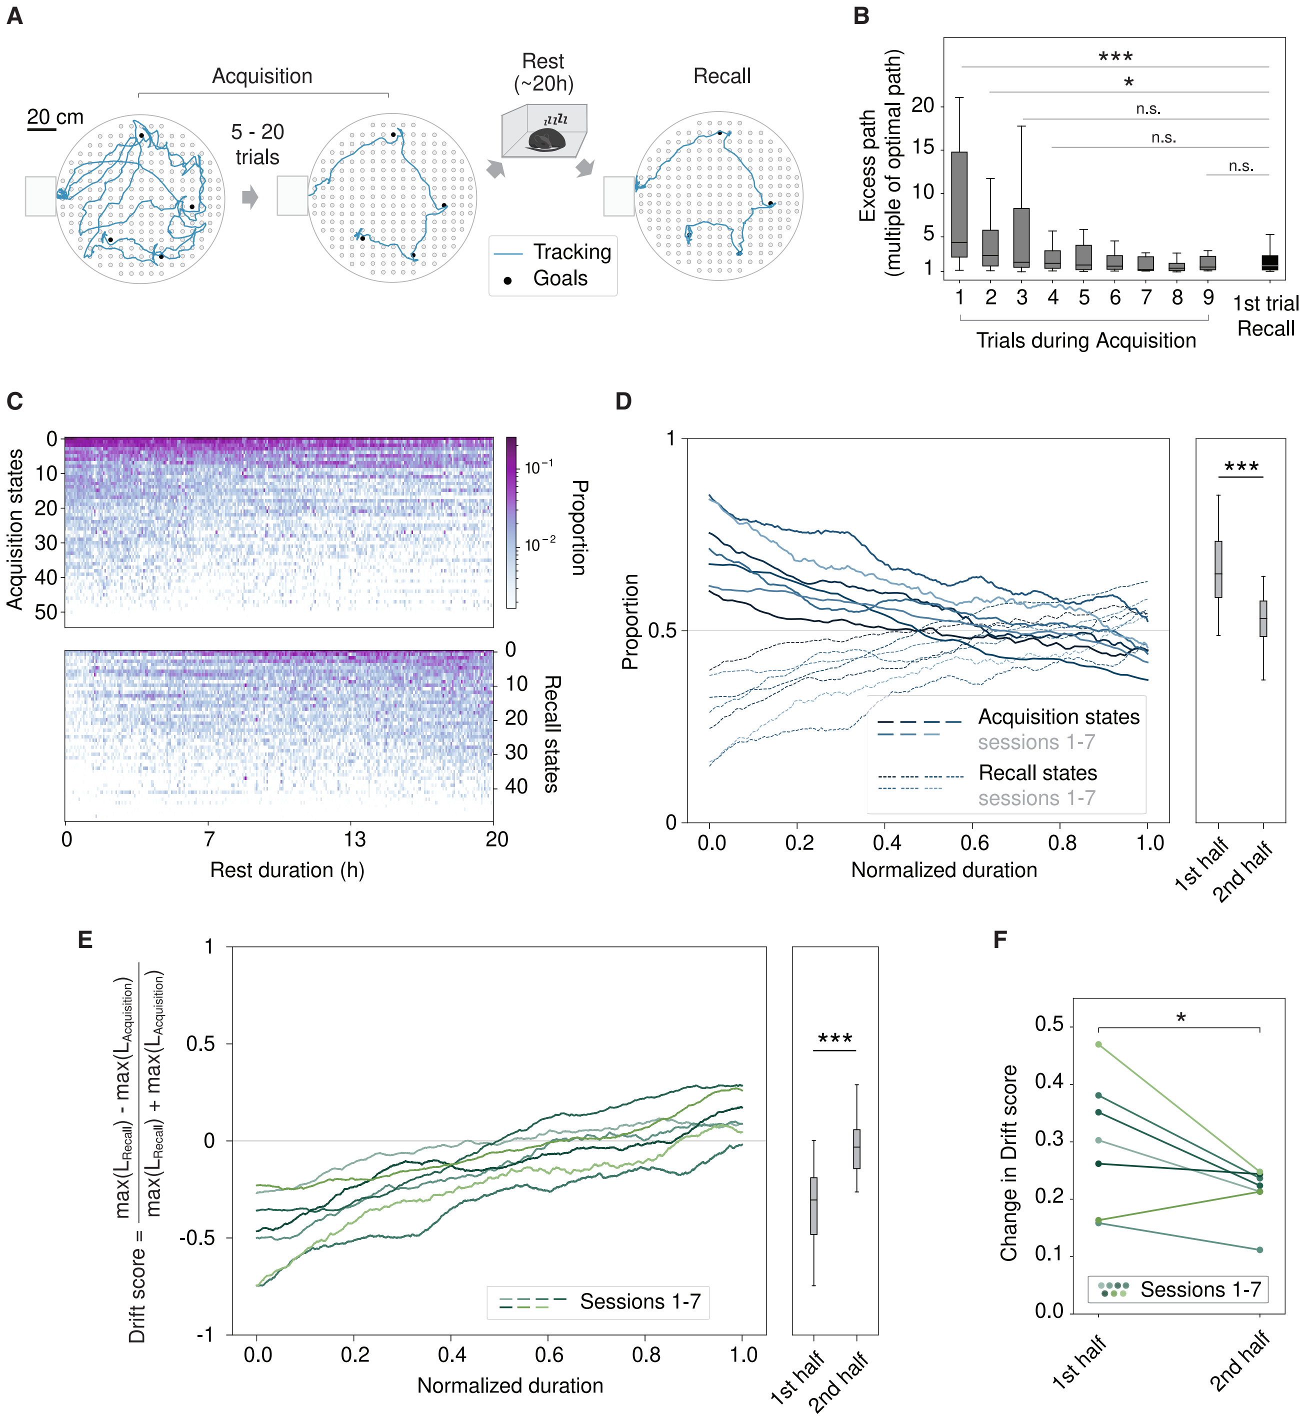

Figure 1. Reactivation drift during prolonged periods of sleep/rest following spatial learning on the cheeseboard

(A) Behavioral sessions illustrated by animal tracking data for one example session: learning a novel set of goal locations (acquisition), long (∼20 h) quiet rest/sleep period (rest), and recall of goal locations (recall).

(B) The excess path is measured as a multiple of the optimal path to retrieve all goals during the first nine trials of memory acquisition and the first trial of recall (*$p < 0.05$, ***$p < 0.001$, n.s. $p > 0.05$, one-sided Mann-Whitney U test, Bonferroni correction).

(C) Instances when acquisition and recall pHMM states were detected, with the color indicating the corresponding proportion that the pHMM state is decoded over a fixed ($n = 600$) time window for one example session.

(legend continued on next page)

32-tetrode microdrives. Continuous recordings were performed for a period of approximately 24 h, covering the learning of a novel set of goal locations on a cheeseboard maze,[9] an interleaved, extended rest period at the home cage, and the recall of the previously learned goal locations (Figure 1A). We extracted the activity of multiple, single units of putative pyramidal cells and interneurons that exhibited stable spike features during the entire recording (Figure S1). Animals learned the novel set of goal locations within the first three trials of the learning session, and they were able to directly retrieve the rewards at subsequent trials, even at the first trial of the recall session the next day (Figure 1B).

We used a Poisson hidden Markov model (pHMM) to identify the activity of different pyramidal cell assemblies during the learning and recall sessions (Figure S2). The model did not use any prior assumption about assembly coding nor the spatial coding of place cells: it was only based on the correlated temporal firing structure of pyramidal population activity in order to identify activation states. To see whether pHMM states expressed spatial selectivity, we plotted the location of the animal at each time a specific state was active. The majority of pHMM states tended to occur only at a narrow set of locations on the cheeseboard maze, many near the goal locations, even though the states were established without prior knowledge of the location of the animal (Figures S3A and S3B). Consequently, the pHMM was able to decode the animal's location with a higher accuracy than that obtained by a standard Bayesian decoding procedure (Figures S3C–S3G), using a cross-validation procedure in which the prediction in each trial was performed by fitting a model using the other trials.

Next, we used the pHMM to trace the reactivation of waking assembly patterns during the prolonged rest periods between the learning and recall sessions. For the majority of the rest period, the animal was immobile and sleeping—or awake but immobile in its home cage. In examining reactivation, we only considered quiet rest periods and excluded periods when the animal moved within the home cage. In the decoding procedure, we used a temporal binning that conserved the number of spikes in any bin (12 spikes, if not stated otherwise). Within the rest period, we differentiated REM sleep periods and the remaining quiet rest NREM periods and identified reactivation events during continuously detected time windows in REM sleep and in SWRs during NREM periods. The incidence of SWRs was initially higher, but, even toward the end of the rest session, their occurrence frequency was >50% of that seen in the beginning (Figure S4E). Reactivations of specific patterns were identified by determining, for each fixed-spike bin, the pHMM state that matched the activity with the highest likelihood, assuming Poisson firing probabilities (Figure S2A; see STAR Methods). To see whether reactivated neuronal patterns underwent a transformation during the prolonged rest periods, we separately compared rest patterns with the set of pHMM states of both acquisition and recall and tracked their likelihoods during rest. We found that in the first half of the rest session, acquisition pHMM states were preferentially decoded (Figures 1C and 1D). In contrast, later acquisition and recall states were decoded with similar probability or the decoding of acquisition states showed a weak preference. This indicated a drift in the reactivation patterns: the majority of activity patterns were initially similar to those expressed in the acquisition stage, but, later, progressively more of them resembled better recall-stage activity patterns seen after the rest session.

Next, we tracked over time how the reactivation of learning-associated neuronal patterns was gradually overtaken by recall-associated patterns. To do so, for each reactivation event, we compared the decoding log-likelihood of the learning and recall pHMM states and a drift score was calculated, defined by the normalized difference of the log-likelihoods (Figure S2A). A negative drift score indicates the reactivation of acquisition pHMM states, whereas a positive score corresponds with patterns more similar to the recall states. The smoothed drift score gradually increased over time, reaching positive values only after the second half of the rest session, after 7–8 h (Figure 1E). However, the speed of change of the drift score was faster during the first half of the rest session than in the second half, suggesting that the majority of new recall states had already started to be active in the first half of the rest session (Figure 1F). We obtained similar results when Bayesian position prediction probabilities were used to calculate the drift score (Figures S3H–S3J) or with the pHMM when $n = 6$ spikes were used (Figures S4F and S4G). We further tested whether similar reactivation patterns could be detected during putative quiet waking periods in which <5 min immobility periods were surrounded by active movement periods (Figures S4A–S4C). We observed similar reactivation drift and the decoding probabilities were somewhat higher in these periods. The amount of active wakefulness did not correlate with the drift score (Figure S4D).

Given that in the second half of the rest session, the proportion of events decoded as acquisition or recall states was similar (40%–60%), we further tested whether these events were not spuriously classified. To this end, we scrambled the cell assignment of spikes by taking 10,000 consecutive spikes (scramble interval <100 s, continuously during REM and consecutive SWRs during NREM) by randomly swapping the cell identity of spike pairs (Figures 2A and 2B). This procedure maintained the temporal patterns of the spikes at the population level: it only scrambled their cell assignment. Consistently, we detected higher decoding probabilities for the original data than for scrambled spike data for both decoded acquisition or recall states. This showed that reactivation events were not random. This difference was not only observed when reactivation patterns during REM and NREM were analyzed separately but also when outside SWR events were tested (Figures S4H–S4K).

(D) The proportion of decoding acquisition over recall pHMM states at different stages of rest for all sessions (i.e., the proportion of acquisition or recall decoding windows identified). The detection probabilities are significantly different in the first and second half of the rest session ($p < 0.001$, two-sided Mann-Whitney U test). Note that the acquisition and recall proportions sum up to one for each session; therefore, session curves are mirror images of each other.

(E) Drift score as a function of normalized rest duration for all sessions. The smoothed ($n = 20,000$) drift scores are different in the first and second half of the rest session ($p < 0.001$, two-sided Mann-Whitney U test).

(F) Drift score changes are stronger during the first half of rest than during the second half of rest ($p < 0.05$, paired t test).

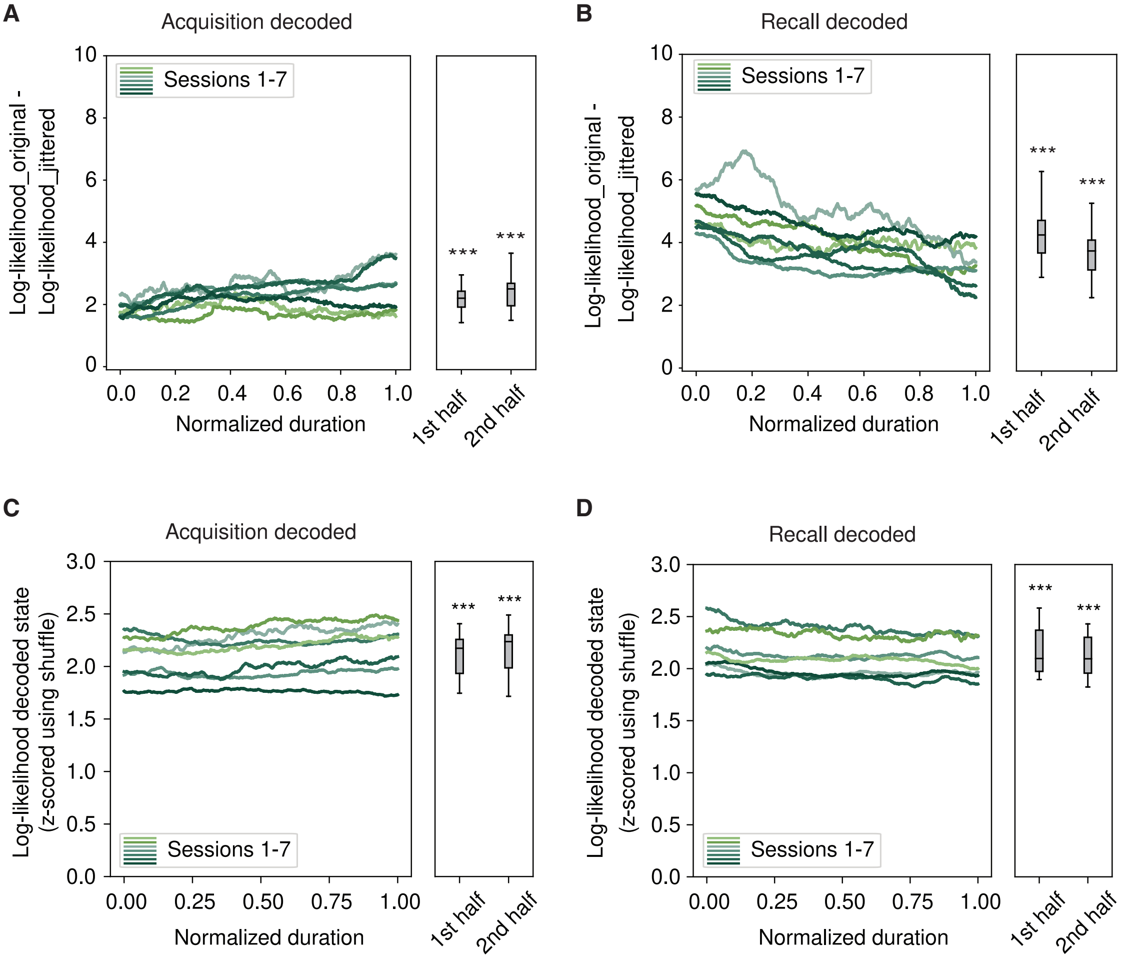

Figure 2. Validation of pHMM-model-based decoding and the influence of rate equalization

(A and B) Difference in log-likelihoods between decoding acquisition (A) and recall states (B) using original data and spike-jittered data. The jittering was performed $n = 100$ times, and the average results are shown. Differences are above zero for the first half and the second half of the rest for all sessions, ***$p < 0.001$ t tests with Bonferroni correction.

(C and D) Log-likelihood of decoded acquisition (C) and recall (D) state, Z scored using shuffle ($n = 100$) in which state vectors of each model are scrambled on a cell-by-cell basis (decoding likelihoods of 100 artificial acquisition and recall states). ***$p < 0.001$, t test with Bonferroni correction.

In addition to the spike jittering, we also performed a pHMM state shuffling procedure in which we scrambled the cells' firing rate across different state vectors of the same model (Figures 2C and 2D). That is, for each cell, we randomly reassigned their firing rates across different pHMM state vectors. We performed this shuffling independently for acquisition and recall models using 100 shuffling iterations. Then, we Z scored the log-likelihood of the original decoded states relative to the shuffled values. This Z score was consistently positive throughout the entire rest session, again supporting the specificity of reactivated patterns during this period.

Next, we further tested whether the decoded model (either acquisition or recall) could be reliably differentiated from the other model. The difference of maximum log-likelihoods of the two models was always >3, indicating at least a 10^3-fold difference in decoding likelihoods (Figures 3A and 3B). In addition, we provided further confirmation as to whether the decoded acquisition and recall states can be statistically separable from the maximum likelihood state of the other model (Figures 3C and 3D). To do so, we took the firing-rate vector of these two (maximum likelihood) states and randomly exchanged the firing rate of the corresponding cells (for each cell, either swapped it or did not) and the log-likelihood differences were calculated while using the scrambled state vectors for the decoding. Scrambling was performed 100 times and the mean and the SD of the log-likelihood differences derived from the scrambled state vectors were calculated. Finally, the original differences (non-scrambled state vectors) were Z scored relative to those from the scrambled vectors. The Z scored values were positive, showing that the likelihood differences with the original state vectors were significantly higher than those of the state scrambled vectors.

In summary, the neural activity during rest gradually changed from being initially better aligned with acquisition patterns to

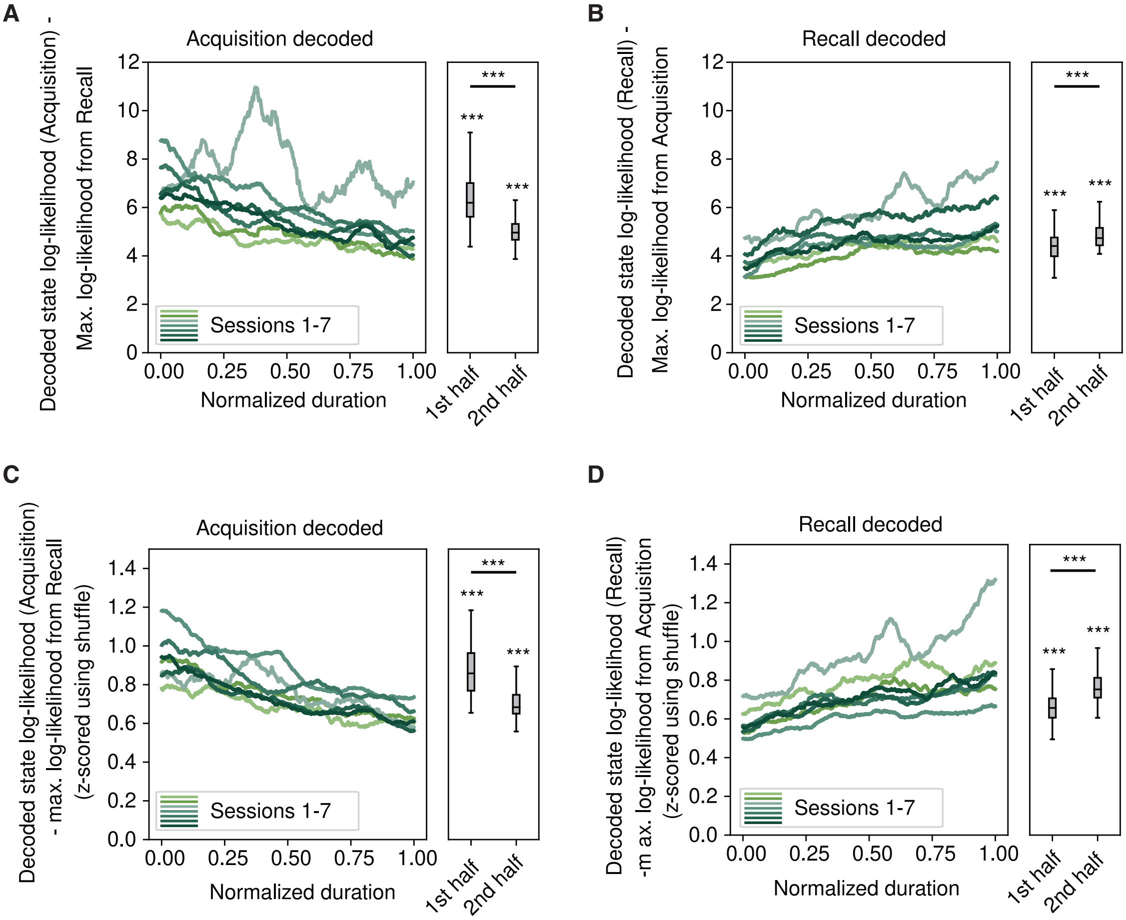

Figure 3. Distinguishing acquisition and recall pHMM states

(A and B) Difference in log-likelihood between the decoded state and the state with the highest likelihood from the other model: (A) using only decoded acquisition states and (B) recall states. Differences are larger than zero for the first and second half of all sessions (all $p < 10e-50$, t test with Bonferroni correction) and they are also significantly different $p < 0.001$, Mann-Whitney U test.

(C and D) *Z* scored difference in log-likelihood between decoded state (acquisition, C or recall, D) and maximum likelihood state from the other model, *Z* scored relative to differences yielded from shuffled cases in which rates of the two state vectors were randomly swapped. The differences were *Z* scored using the mean and SD of the difference using 100 artificially generated states for the decoding in which the firing rate of the max likelihood states of the two models was randomly swapped for each cell. Values for the first and second half are larger than 0 for all sessions (all $p < 10e-50$, t test with Bonferroni correction), and they are also significantly different $p < 0.001$, Mann-Whitney U test.

being increasingly similar to the recall-associated neural activity at the end of rest.

We next examined whether reactivation drift exhibited differences in REM and NREM periods. When we plotted drift scores over time after performing a smoothing window averaging, we observed temporal fluctuations with a zigzagging appearance. We asked how these repeated upward and downward trends aligned with the transitions between the NREM and REM phases. We found that the smoothed drift scores tended to increase during NREM while decreasing during REM periods, hence demonstrating that the fluctuations we saw were driven by the NREM-REM cycles (Figure 4A). To quantify this effect, we calculated the change in drift scores for each REM and NREM epoch at a fine temporal scale ($n = 20$ smoothing) to limit instances in which a smoothing window contained mixed states (Figure S4L). The average drift score changes were positive for NREM and negative for REM, with only about a quarter of the REM and NREM epochs showing the opposite tendencies (Figure 4B, i.e., positive REM and negative NREM). Consequently, the cumulative effect of NREM periods on the drift score was positive, whereas REM periods showed a negative cumulative effect (Figures 4C and S3J). We found a difference of two orders of magnitude when we compared the cumulative change in the drift score throughout the rest with the epoch-by-epoch change in NREM and REM (Figure 4D). Short- and longer-term changes in drift score occurring during NREM and REM epochs were not correlated (Figure S3K). Hence, the drift score fluctuated strongly on shorter timescales without a clear direction, whereas the long-term change was directed and much smaller in magnitude, which implies that different processes drive the short-term fluctuation and the long-term drift.

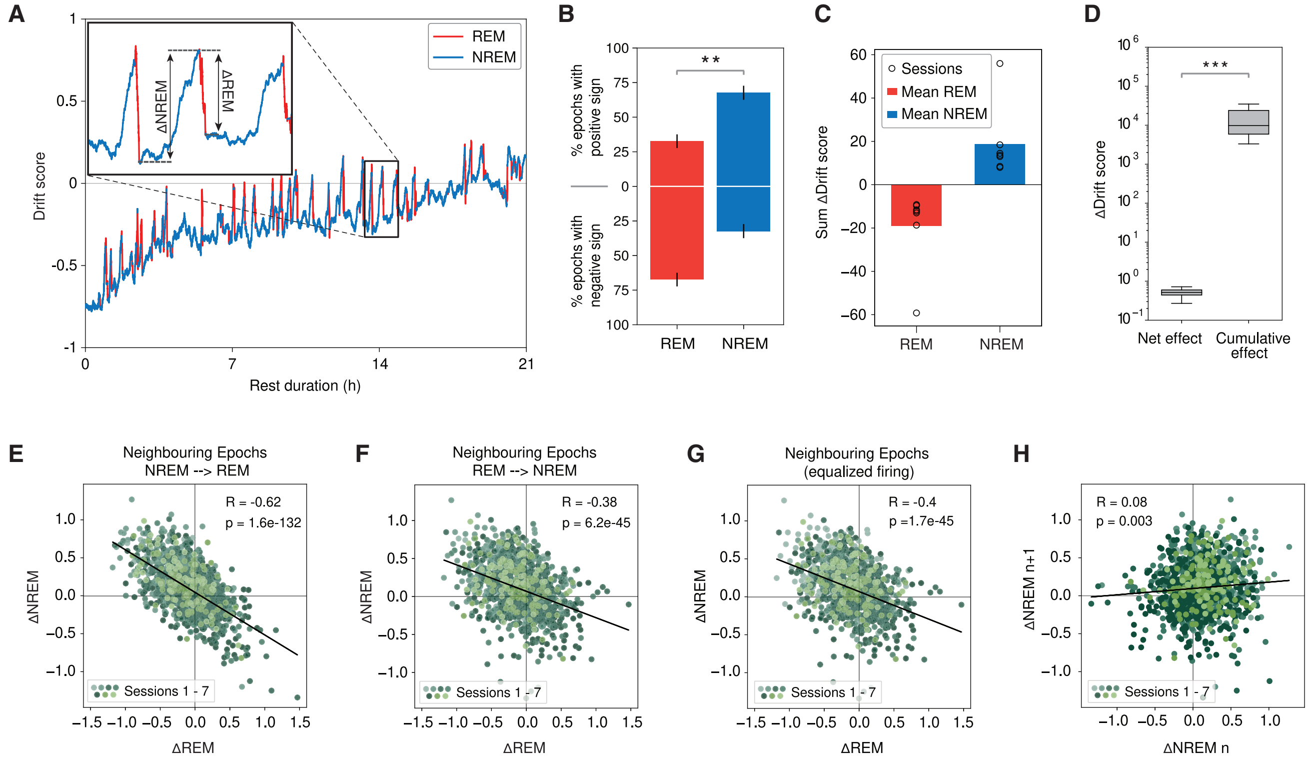

Figure 4. NREM periods accelerated the reactivation drift, whereas REM periods countered it

(A) Smoothed ($n = 900$) drift score as a function of rest duration (REM periods in red, NREM periods in blue) for one example session. Inset: time-expanded trace illustrating the ΔDrift score calculations.

(B) Contributions of REM and NREM epochs to drift. For each epoch, the ΔDrift score was computed as the difference in the drift score between the end and beginning using $n = 20$ window smoothing. The percentage of epochs with positive and negative values are depicted for REM (red) and NREM (blue). Data from all sessions (mean ± SEM, $^{**}p < 0.01$, two-sided Mann-Whitney U test).

(C) Summed ΔDrift score for REM and NREM epochs for all sessions ($p < 0.01$, two-sided Mann-Whitney U test).

(D) The net effect of change in drift score (difference in drift score between the beginning and the end of rest) and the cumulative effect of change in drift score (sum of absolute ΔDrift score for REM and NREM periods). Data from all sessions ($^{***}p < 0.001$, two-sided Mann-Whitney U test).

(E) ΔDrift score values of NREM with subsequent REM epochs ($R = -0.62$, $p = 1.6e-132$).

(F) ΔDrift score values of REM with subsequent NREM epochs ($R = -0.38$, $p = 6.2e-45$).

(G) ΔDrift score values for NREM and REM epochs after equalizing firing rates across the neighboring epochs ($R = -0.4$, $p = 1.7e-45$).

(H) ΔDrift score values for subsequent NREM epochs ($R = 0.08$, $p = 0.003$).

Next, we tested whether the magnitude of up- and downward shifts in the drift scores across NREM and REM periods were independent or whether the drift in subsequent epochs was interrelated. The drift score changes in neighboring NREM vs. REM periods were negatively correlated, and the NREM to REM drift score changes exhibited stronger correlations than the REM to NREM ones ($p = 3e–15$, Fisher's z test, Figures 4E and 4F). However, the correlations of the drift score changes were weaker when subsequent NREM periods were compared (Figure 4H) and not significant for subsequent REM periods ($R = 0.01$, $p = 0.85$). Overall, these results suggested that NREM periods drove the drift of reactivated assemblies toward the ones that will re-emerge during the recall session, whereas REM periods countered this effect. The correlations also suggest that the NREM drift influences the magnitude of subsequent REM drift more than REM drift influences the subsequent NREM changes.

To test the contribution of individual cells to the observed reactivation drift, we characterized the firing-rate changes of neurons between the acquisition and the recall period (Figure 5A; STAR Methods). We observed that some pyramidal cells exhibited stable firing rates across the two sessions whereas others either increased or decreased their firing rate from one to the other (Figure 5B).

When we looked at the mean rate of the cells as specified in their pHMM mean-firing-rate vectors (Figure S2A), the same trend appeared when acquisition and recall states were compared (Figure S5A). This showed that, with the activation of recall pHMM patterns later on during the rest session, cells showed more similar firing-rate trends to those seen later during the recall session. More than half of the recorded cells were from the decreasing group, and only a smaller portion of cells had stable or increasing firing rates over the entire experiment (Figure S5B). Cells with persistent activity, on average, exhibited lower firing rates than the rate-changing groups but maintained a comparatively higher rate during the rest session (Figures 5C and S5C–S5E). The burstiness of the rate-changing

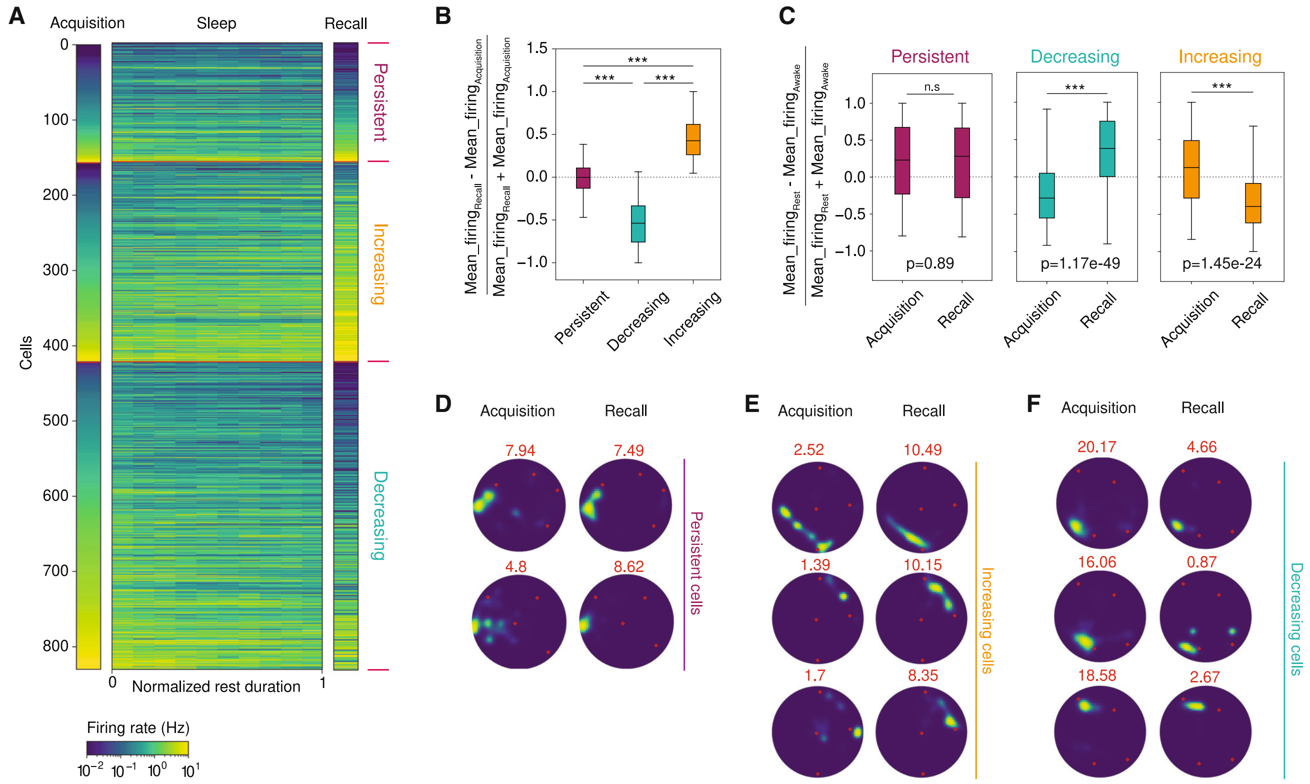

Figure 5. Pyramidal cell classification on the basis of firing-rate change

(A) Average firing rates of persistent, increasing, and decreasing cells during acquisition, recall, and fixed periods of rest. Cells from all sessions are shown. Cells were classified depending on whether their firing-rate distributions exhibited significant differences ($p < 0.05$) between acquisition and recall.

(B) Normalized firing-rate change from acquisition to recall for persistent, decreasing, and increasing cells (all $p < 0.001$, two-sided Mann-Whitney U test, Bonferroni correction).

(C) Normalized firing-rate change between rest vs. acquisition or recall (n.s. $p > 0.05$, ***$p < 0.001$, two-sided Mann-Whitney U test).

(D–F) Example rate maps for persistent (D), increasing (E), and decreasing (F) cells. Top numbers: peak firing rate in Hz, red dots: reward locations. Blue marks the low-firing-rate region and yellow the higher-rate place field region.

groups was different, but they tended to exhibit higher burst propensity when their rate was lower (Figures S6E and S6F). We did not find differences in SWR firing gain or waveform stability between the different subsets (Figures S5F and S5H). Different cells recorded from the same electrode could exhibit different rate modulations (i.e., belong to different groups) (Figure S5G) and, accordingly, the polarity of the recorded sharp waves from the electrode that a cell was recorded from did not differentiate cell groups (Figures S6G and S6H). The difference between the preferred theta phase in REM and waking theta (measured during acquisition) was also similar (Figures S6A–S6D). Spatial information of the persistent cells was lower compared with decreasing cells during the acquisition and lower compared with the increasing cells during recall (Figure S5I). Nevertheless, we did not observe major differences in the goal-related remapping during learning or recall between persistent and rate-changing cells (Figures S7A–S7D).

When we performed the reactivation drift analysis using only the persistent cells, the drift score remained stable near zero across the entire rest period, indicating a high similarity of acquisition and recall patterns and their consistent reactivation during sleep (Figure 6). However, reactivations by the other two groups yielded drift score changes similar to those calculated with all cells. This result suggests that firing-rate changes of the rate-changing cell groups were the primary drivers of the reactivation drift. Furthermore, we also found that the spatial coding of the environment remained more similar for the persistent cell group than for the rate-changing groups, as measured with population vector similarity across acquisition and recall (Figures S5I and S5J).

The fact that only rate-changing cells contributed to the reactivation drift might suggest that solely the excitability changes drove the reactivation drift. However, this was not the case. To show this, we equalized the firing rate of the cells throughout the entire rest session by dividing the rest periods into five segments and randomly down-sampled the activity of each cell so that their average rate in each segment corresponded to that of the lowest rate segment in the original data (Figure S4M). In this instance, we observed a similar reactivation drift to the original case (Figure S4N), showing that rate changes themselves were not necessary for the emergence of reactivation drift.

To further investigate the contribution of pyramidal rate changes to reactivation drift, we tested whether rate changes seen during individual REM and NREM epochs correlated with

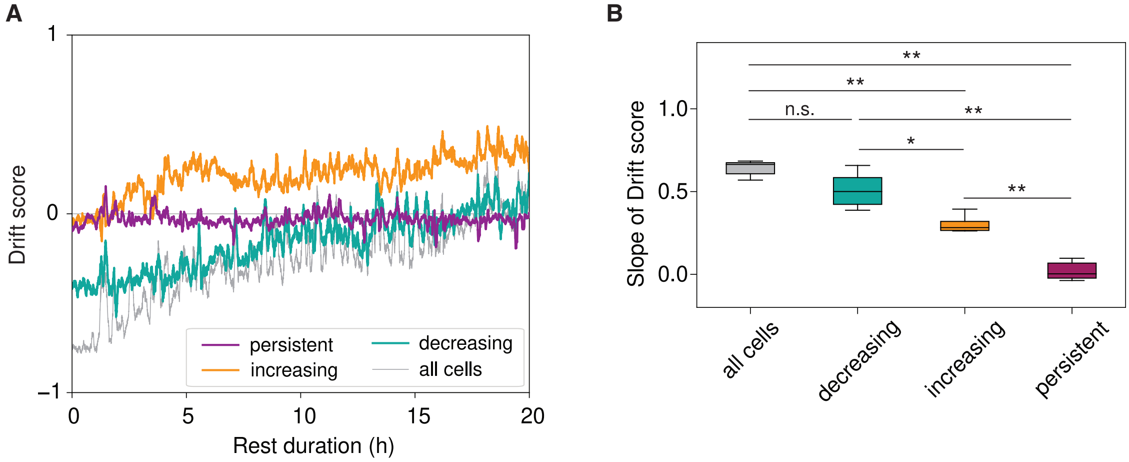

Figure 6. Contribution of different firing-rate modulation cell groups to reactivation drift

(A) Drift score using either all cells (gray), only persistent cells (violet), only increasing cells (orange), or only decreasing cells (turquoise) for one example session.

(B) Slope fit to drift score (**$p < 0.01$, *$p < 0.05$, n.s. $p > 0.05$, two-sided Mann-Whitney U test, Bonferroni correction).

the magnitude of drift score changes at the same epochs: for the increasing group, we detected a positive correlation, whereas a negative correlation was seen for the decreasing cells (Figures 7A–7G). Similar to the drift score changes, neighboring REM and NREM periods exhibited correlated firing-rate changes for both the decreasing and increasing groups, with a stronger effect for NREM periods being followed by REM periods ($p = 1.23\text{e}{-22}$ for decreasing cells, $p = 4.91\text{e}{-19}$ for increasing cells, Fisher's z test, Figures S7E–S7H). Firing rate changes between rest epochs of the same type were only weakly correlated (Figures S7I–S7L).

Previous work suggested that, in NREM sleep, overall network excitation increases, which is countered by rate reduction in the subsequent REM period.[18,19] Therefore, we checked whether a non-specific change in network excitability alone could drive the REM/NREM reactivation fluctuation. To test this, we equalized the population rate in subsequent NREM-REM periods by randomly omitting spikes. Yet, even in this case, the changes in drift score in subsequent NREM and REM epochs were negatively correlated (Figure 4G). Moreover, macroscopic measures of network activity, such as Z scored pyramidal network firing rate or LFP amplitude, did not correlate with the change of drift score during REM or NREM periods (all $r < 0.06$, all $p > 0.05$)—nor did the number of SWR during NREM correlate with it ($r = -0.01$, $p = 0.85$).

Our data showed that rate-increasing cells preferably fired with specific recall assemblies (pHMM states), whereas rate-decreasing cells preferred acquisition assemblies (Figures S7M–S7T). Therefore, we further checked whether the transient rate increases of the increasing group during NREM and the similar transient REM increases of the decreasing group are assembly specific (i.e., pHMM-state related). We reasoned that, in the case of an assembly dependent rate increase of these cells, the decoding likelihoods should also increase. However, the decoding likelihood should decrease if the cells' rate increases were independent from the activity of their assembly partners. Indeed, the increasing group, which is preferentially active when the recall state is decoded (Figure S5A), exhibited a positive correlation with the change recall model maximum log-likelihood changes (Figures S7P and S7T). By contrast, the decreasing group, active when the acquisition state is decoded, exhibited a similar correlation with acquisition maximum log-likelihood changes (Figures S7M and S7Q).

Next, we tested whether the interneuron activity was related to reactivation drift. The activity of some interneurons strongly reflected the drift score changes in REM and NREM epochs (Figures 8A and 8B). In agreement with this, 25.2% of the interneurons exhibited a significant correlation between their firing rate and the change in drift score (Figure 8C). Some of these interneurons increased their firing rates with the magnitude of rate changes during both REM and NREM epochs, whereas others reduced it.

After separating REM and NREM epochs, a significant fraction of interneurons only maintained a relationship between mean firing and the change in drift score for REM epochs (Figures S8A–S8N). However, interneurons exhibited a reversed relationship in REM and NREM periods between mean firing rate and change in drift score: if an interneuron's firing rate showed a positive relationship with the change in drift score in REM, the relationship was negative in NREM and vice versa (Figure 8D). This effect modulated all recorded interneurons regardless of the strength of the firing rate to drift relationship. In summary, the firing rates of a subset of interneurons mirrored the fluctuations of the drift score. Moreover, the whole interneuron population exhibited a reversed coupling to REM and NREM periods in terms of their net firing-rate modulation, with some exhibiting preferred rate increases either in REM or NREM compared with the other state.

DISCUSSION

Here, we demonstrated that reactivated patterns in the hippocampus reorganize during prolonged rest periods (17–20 h) after spatial learning of novel goal locations in a familiar environment.

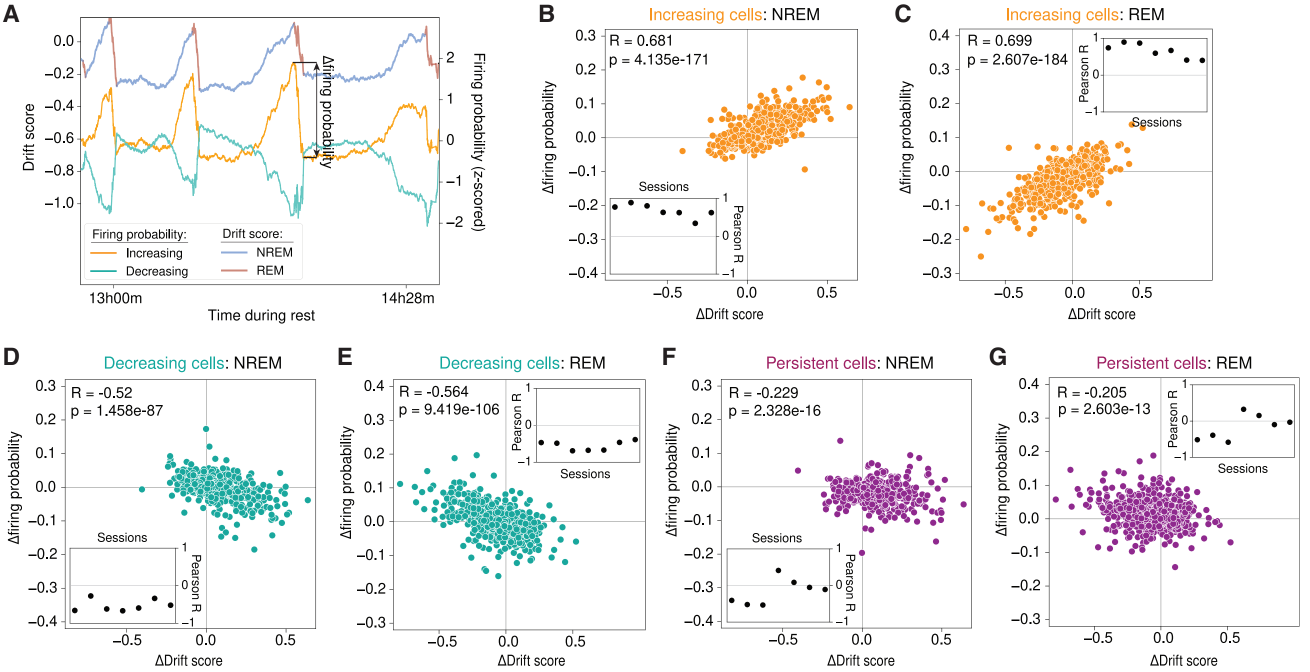

Figure 7. Reactivation drift correlates with the firing-rate modulations of pyramidal populations

(A) Illustration of drift score and firing probability fluctuations of decreasing and increasing cells for one example session.

(B–G) Correlation between ΔDrift score and the change in firing probability for increasing cells during NREM and REM periods. Correlation coefficients (R) and corresponding p values are shown on the top left. Inset: R values per session.

In the first 8–10 h of rest, these reactivated patterns primarily represented the previous learning session patterns, but, later, they were gradually overtaken by those resembling future patterns expressed in the subsequent recall session. We also found that this reactivation drift was antagonistically modulated by the succession of NREM and REM epochs: whereas NREM periods favored the drift toward new patterns of the recall session, REM periods had the opposite effect. Hence, we observed the gradual reorganization of reactivated hippocampal assembly patterns, and the reorganization was regulated by NREM-REM cycles.

Reactivation drift reflects place map reorganization

Our ability to detect the progressive transformation of hippocampal reactivated patterns relied on their continuous monitoring over prolonged periods of sleep/rest. Previous studies in which intermediate sleep durations (<6 h) were used observed a relatively stable reactivation expression.[16,17] In our case, the majority of the novel recall patterns emerged in the first half of the rest periods, and reactivation drift occurred primarily in this period, during which new patterns gradually replaced the prior learning patterns. However, after about 8–10 h, recall assembly patterns matched or overtook the learning patterns. Therefore, the more frequent occurrence of the old patterns during the earlier part of the sleep could explain why previous studies of shorter sleep durations observed relatively stable reactivation even after the prior exploration of novel environments.

However, there are experimental data suggesting that reactivated sleep patterns do not exclusively represent recent experiences. Even early reactivation studies were able to detect a weak but significant similarity between sleep activity patterns and subsequent waking patterns of familiar environments under conditions in which the animal visited the familiar environment a day before at the earliest.[20] This suggests that less recent experiences are also reactivated in sleep, including those the animal experienced in previous days. However, later preplay was observed in which sleep temporal patterns were similar to place-cell firing patterns that emerge later in a novel environment.[21–23] Even in familiar environments, one can observe a representational drift of place maps, where some new place cells emerge while others disappear in later exploration of the same environment.[24–26] Similar to preplay, neuronal patterns during sleep can also reflect activity patterns of place cells that newly emerge in a familiar environment.[27] These data all suggest that, after prolonged sleep, some of the patterns can reflect novel patterns seen after sleep in novel or familiar environments.

Can the reactivation drift in our data be related to preplay or the representational drift of place representations seen across familiar environments? Our data analysis used a pHMM model to identify distinct cell assemblies, which only took into account the correlated temporal firing structure of the recorded pyramidal cell population. However, the waking activation of pHHM states exhibited spatial selectivity, indicating that these corresponded to distinct assemblies of place cells representing different locations. Accordingly, the activation of recall states in the prior rest session suggests that some of the updated place cell assemblies of the recall session already emerged in prior sleep. We could further confirm this effect by using a Bayesian place-decoding reactivation analysis (Figures S3H–S3J).

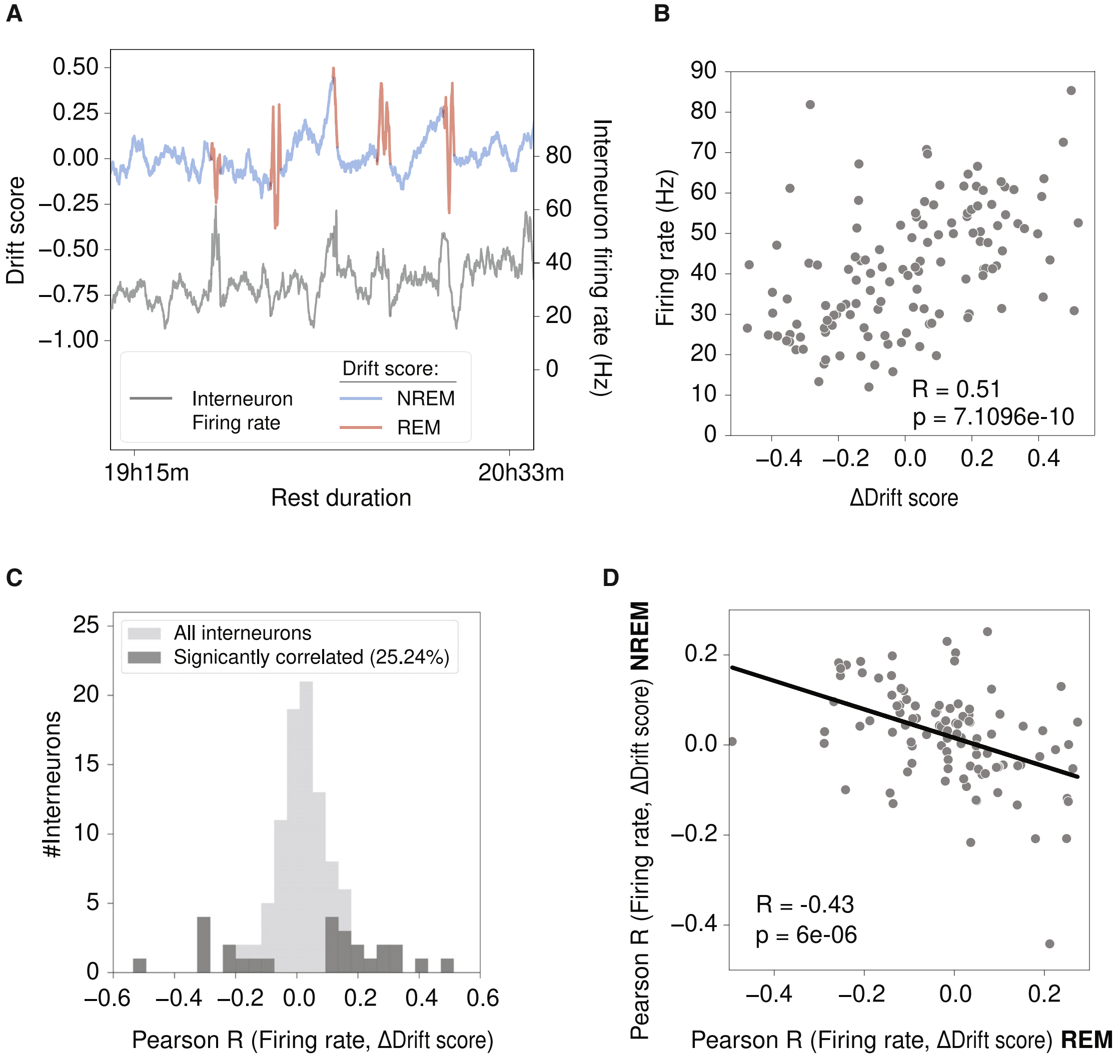

Figure 8. For a subset of interneurons, change in drift score correlates with interneuron within epoch firing rate

(A) Example of the correlated fluctuation of the drift score and the firing rate of an interneuron.

(B) Correlation of firing rate and ΔDrift score for the interneuron from (A) (R = 0.51, $p = 7.11e-10$).

(C) Distribution of Pearson R values for interneuron firing rate and delta drift score using REM and NREM epochs. A significant proportion of interneurons exhibited a significant ($p < 0.05$) correlation between firing rate and ΔDrift score ($p < 0.001$, binomial test).

(D) Pearson R values for interneuron firing rate and ΔDrift score were calculated separately in NREM and REM epochs (R = −0.43, $p = 6e-6$).

Considering that we observed this effect in a familiar environment (albeit during novel spatial goal learning), our findings may also be related to the representational drift observed across hippocampal place cells over repeated exposure to familiar environments.[24,25,28] However, our experimental conditions were different from the preplay studies in which an entirely novel environment was introduced. In our case, the animal learned a novel set of reward locations in a familiar environment and the prolonged rest period separated the next day's recall. This condition was actually similar to a recent work in which extra-stable recording in two-photon imaging was performed daily while animals learned and recalled a reward zone on a running belt.[29] In this study, 35% of place cells maintained similar spatial representations across days. This proportion is somewhat similar (25%) to our stable cell groups that maintained stable firing rates across learning and recall. The spatial population vectors of the stable cells strongly correlated (r = 0.9) across learning and recall sessions, demonstrating that, indeed, this group maintained stable spatial representations for longer than a day. However, the other two groups exhibited lower but still relatively strong population vector correlations (decreasing 0.5 and increasing 0.7), indicating that some of these cells also maintained similar place fields while others cells may have remapped. Also, a significant portion of the rate-changing cells remained active during the entire recording (top 50% exhibited >0.3 Hz). However, the relatively higher population vector correlations and sustained activity at a lower firing rate suggest that many of the rate-changing cells may have undergone rate remapping.

Stable and rate-changing pyramidal cell groups

Several studies demonstrated that hippocampal pyramidal cells are not uniform and can be subdivided into groups based on their anatomical or physiological properties. Initially, CA1 pyramidal cells have been differentiated according to their calbindin expression.[30–35] Later work showed that CA1 pyramidal cells could be subdivided according to their oscillatory firing properties, firing rate, and burst propensities.[36] In our classification, we compared the firing-rate distribution of cells across learning and recall sessions to separate them into stable-rate persistent cells and rate-changing groups that either increased or decreased their rate from learning to recall. We showed that persistent cells alone did not exhibit reactivation drift, whereas the rate-changing groups did. Interestingly, about half of the pyramidal cells belonged to the rate-reducing group in which cells also exhibited a rate reduction whenever the newly emerging recall pHMM states were activated at rest. This is consistent with the homeostatic rate reduction role of sleep and further suggests a more sparse place representation to emerge after sleep.[37] We also observed a negative correlation between firing-rate changes across REM and NREM sleep, which is in line with the opposing role of these states in firing-rate regulation.[18]

In relation to replay and preplay, plastic and rigid cell groups have been differentiated: rigid cells participated in the preplay of firing sequences, whereas plastic cells refined replay sequences in subsequent sleep without contributing to preplay.[22] Overall, in that study, the rigid cells exhibited higher firing rates and reduced spatial selectivity compared to the plastic group. As mentioned above, reminiscent of preplay, we observed the emergence of updated recall assembly patterns before and during the rest session. However, conditions in our case were different; the rate-changing cells exhibited more similar firing characteristics to the rigid cell groups seen in preplay: they had overall higher firing rates than the persistent cells.

Network mechanism of reactivation drift

As we saw above, the persistent cells exhibited stable reactivation and spatial coding for the entire duration of the recordings. These cells probably overlap with those place cells that, in calcium imaging experiments, maintained stable place fields across days.[29,38] These cells can enable continuity and provide a reference frame for the transformation of reactivation patterns. At the same time, we expect that the update of reactivated spatial assemblies would require plasticity that enables new members of the increasing group to associate with existing members of the persistent group and the decreasing group to decouple from them.

SWRs themselves may provide the means for plastic changes to enhance the assembly association of new members.[39] Previously, it was shown that if a CA1 cell is repeatedly excited during a SWR, that cell will increase its SWR-associated firing after the pairing.[40] We expect that assembly specific SWR pairing will increase the cell's firing only when that paired assembly is activated during SWRs, considering that SWR-like synchronization enhances the plasticity of cells that are co-active.[41,42] The drift toward the updated assemblies primarily occurred during NREM periods in the presence of SWRs. Therefore, we speculate that the association of new members to an assembly may be initiated by a random process in which some non-member cells may start to fire with a specific assembly. If that cell has a lower threshold for plasticity, even a few random pairing events may initiate a self-reinforcing process in which pairing will lead to increased activation probability with an assembly, which will lead to further pairing and a further increase in assembly association.

In parallel with the opposing role of REM sleep in reactivation drift, plastic changes during REM sleep may favor assembly stabilization instead of assembly drift. Plastic changes that are thought to enhance the stabilization of newly formed place fields in waking theta oscillations may operate in a similar manner during REM sleep as well. These may also enable the strengthening of existing assemblies. This process may also extend to updated assemblies that emerge in previous NREM periods and enable them to be maintained in subsequent waking activity as well. Given that theta oscillations can facilitate both long-term potentiation and long-term depression,[43,44] not only could associations of new members be enhanced during REM sleep but the decoupling of other cells from the rate-decreasing group could also occur.

A number of additional factors may also impact different reactivation dynamics and associated circuit functions during REM and NREM epochs.[45] Many of these could be attributed to different physiological states due to the different levels of neuromodulation provided by acetylcholine or other non-specific neurotransmitters, including serotonin.[46,47] Moreover, sleep stages are also different in terms of large-scale network interactions across the entire brain.[6,47–49] One could speculate that enhanced out-of-assembly activation during SWR events might have originated from bi-directional interactions of the hippocampus with other cortical centers.[50] In the same way, the strengthening of existing assemblies during REM periods could be allowed by the relative isolation of the hippocampus from external sources of information, which is thought to be a hallmark of REM periods of sleep.[51–53]

The process we outlined may point to a general sleep mechanism that is not restricted to the case of updated spatial maps in familiar environments or similar spatial learning contexts seen across days. Similar processes may help the rapid initial formation of novel cognitive maps[54] as well and may even enable preplay.[21–23] Ultimately, updated assemblies that emerge during prolonged sleep/rest periods could be used to provide the initial building block for cognitive maps that emerge during the first exposure to a novel environment. Unlike fixed assemblies, this can enable a wide variety of assembly patterns to emerge in sleep, ensuring the emergence of diverse assemblies representing distinct cognitive maps. Similar assembly reorganization involving the gradual morphing of assemblies in sleep could also take place in other brain areas.

RESOURCE AVAILABILITY

Lead contact

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

Materials availability

This study did not generate new unique reagents.

Data and code availability

- Data reported in this paper will be shared by the lead contact upon request.
- All original code has been deposited at Zenodo and is publicly available at https://doi.org/10.5281/zenodo.14803066 as of the date of publication.
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

ACKNOWLEDGMENTS

We thank Andrea Cumpelik, Lisa Genzel, and Freya Ólafsdóttir for comments on an earlier version of the manuscript. This work was supported by the European Research Council (281511) and Austrian Science Fund (FWF I3713).

AUTHOR CONTRIBUTIONS

F.S., P.B., and J.C. designed the study. P.B. conducted the experiments and P.B., F.S., and L.B. analyzed the data. F.S., L.B., and J.C. wrote the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

STAR★METHODS

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

- KEY RESOURCES TABLE
- EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS
- METHOD DETAILS
    - Surgery
    - Data Acquisition, Training and Behavior
- QUANTIFICATION AND STATISTICAL ANALYSIS
    - Spike Detection, Sorting and Stability
    - Sleep classification
    - Sharp wave ripple detection
    - Quantification of stability of clustering features over time
    - Cell separation over time
    - Clustering feature stability
    - Measuring excess path during acquisition and recall
    - Poisson Hidden Markov Model (pHMM) and model fitting
    - Rate map generation for Bayesian decoding
    - Decoding rest activity using pHMM
    - Decoding rest activity using Bayesian decoding
    - Quantifying drift during rest
    - First vs. second half of rest
    - Net effect vs. cumulative effect of drift
    - Opposing effects on drift in NREM and REM
    - Calculation of firing rate, LFP amplitude, and SWR number per REM/NREM epoch
    - Firing probability changes in NREM and REM

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.neuron.2025.02.025.

Received: October 18, 2023
Revised: August 1, 2024
Accepted: February 25, 2025
Published: March 24, 2025

REFERENCES

1. Diekelmann, S., and Born, J. (2010). The memory function of sleep. Nat. Rev. Neurosci. *11*, 114–126. https://doi.org/10.1038/nrn2762.
2. Klinzing, J.G., Niethard, N., and Born, J. (2019). Mechanisms of systems memory consolidation during sleep. Nat. Neurosci. *22*, 1598–1610. https://doi.org/10.1038/s41593-019-0467-3.
3. Buzsáki, G. (1989). Two-stage model of memory trace formation: a role for "noisy" brain states. Neuroscience *31*, 551–570. https://doi.org/10.1016/0306-4522(89)90423-5.
4. Wilson, M.A., and McNaughton, B.L. (1994). Reactivation of hippocampal ensemble memories during sleep. Science *265*, 676–679. https://doi.org/10.1126/science.8036517.
5. Pavlides, C., and Winson, J. (1989). Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. J. Neurosci. *9*, 2907–2918. https://doi.org/10.1523/JNEUROSCI.09-08-02907.1989.
6. Logothetis, N.K., Eschenko, O., Murayama, Y., Augath, M., Steudel, T., Evrard, H.C., Besserve, M., and Oeltermann, A. (2012). Hippocampal–cortical interaction during periods of subcortical silence. Nature *491*, 547–553. https://doi.org/10.1038/nature11618.
7. McClelland, J.L., McNaughton, B.L., and O'Reilly, R.C. (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. https://doi.org/10.1037/0033-295X.102.3.419.
8. Brodt, S., Inostroza, M., Niethard, N., and Born, J. (2023). Sleep—A brain-state serving systems memory consolidation. Neuron *111*, 1050–1075. https://doi.org/10.1016/j.neuron.2023.03.005.
9. Dupret, D., O'Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. Nat. Neurosci. *13*, 995–1002. https://doi.org/10.1038/nn.2599.
10. Grosmark, A.D., Sparks, F.T., Davis, M.J., and Losonczy, A. (2021). Reactivation predicts the consolidation of unbiased long-term cognitive maps. Nat. Neurosci. *24*, 1574–1585. https://doi.org/10.1038/s41593-021-00920-7.
11. O'Neill, J., Senior, T.J., Allen, K., Huxter, J.R., and Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. Nat. Neurosci. *11*, 209–215. https://doi.org/10.1038/nn2037.
12. Cheng, S., and Frank, L.M. (2008). New experiences enhance coordinated neural activity in the hippocampus. Neuron *57*, 303–313. https://doi.org/10.1016/j.neuron.2007.11.035.
13. Girardeau, G., Benchenane, K., Wiener, S.I., Buzsáki, G., and Zugaro, M.B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. Nat. Neurosci. *12*, 1222–1223. https://doi.org/10.1038/nn.2384.
14. Gridchyn, I., Schoenenberger, P., O'Neill, J., and Csicsvari, J. (2020). Assembly-Specific Disruption of Hippocampal Replay Leads to Selective Memory Deficit. Neuron *106*, 291–300.e6. https://doi.org/10.1016/j.neuron.2020.01.021.
15. Sawangjit, A., Oyanedel, C.N., Niethard, N., Salazar, C., Born, J., and Inostroza, M. (2018). The hippocampus is crucial for forming non-hippocampal long-term memory during sleep. Nature *564*, 109–113. https://doi.org/10.1038/s41586-018-0716-8.
16. Giri, B., Miyawaki, H., Mizuseki, K., Cheng, S., and Diba, K. (2019). Hippocampal Reactivation Extends for Several Hours Following Novel Experience. J. Neurosci. *39*, 866–875. https://doi.org/10.1523/JNEUROSCI.1950-18.2018.
17. Hwaun, E., and Colgin, L.L. (2019). CA3 place cells that represent a novel waking experience are preferentially reactivated during sharp wave-ripples in subsequent sleep. Hippocampus *29*, 921–938. https://doi.org/10.1002/hipo.23090.

18. Grosmark, A.D., Mizuseki, K., Pastalkova, E., Diba, K., and Buzsáki, G. (2012). REM Sleep Reorganizes Hippocampal Excitability. Neuron *75*, 1001–1007. https://doi.org/10.1016/j.neuron.2012.08.015.
19. Miyawaki, H., and Diba, K. (2016). Regulation of Hippocampal Firing by Network Oscillations during Sleep. Curr. Biol. *26*, 893–902. https://doi.org/10.1016/j.cub.2016.02.024.
20. Kudrimoti, H.S., Barnes, C.A., and McNaughton, B.L. (1999). Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and EEG dynamics. J. Neurosci. *19*, 4090–4101. https://doi.org/10.1523/JNEUROSCI.19-10-04090.1999.
21. Dragoi, G., and Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. Nature *469*, 397–401. https://doi.org/10.1038/nature09633.
22. Grosmark, A.D., and Buzsáki, G. (2016). Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. Science *351*, 1440–1443. https://doi.org/10.1126/science.aad1935.
23. Farooq, U., Sibille, J., Liu, K., and Dragoi, G. (2019). Strengthened Temporal Coordination within Pre-existing Sequential Cell Assemblies Supports Trajectory Replay. Neuron *103*, 719–733.e7. https://doi.org/10.1016/j.neuron.2019.05.040.
24. Geva, N., Deitch, D., Rubin, A., and Ziv, Y. (2023). Time and experience differentially affect distinct aspects of hippocampal representational drift. Neuron *111*, 2357–2366.e5. https://doi.org/10.1016/j.neuron.2023.05.005.
25. Khatib, D., Ratzon, A., Sellevoll, M., Barak, O., Morris, G., and Derdikman, D. (2023). Active experience, not time, determines within-day representational drift in dorsal CA1. Neuron *111*, 2348–2356.e5. https://doi.org/10.1016/j.neuron.2023.05.014.
26. Mankin, E.A., Diehl, G.W., Sparks, F.T., Leutgeb, S., and Leutgeb, J.K. (2015). Hippocampal CA2 Activity Patterns Change over Time to a Larger Extent than between Spatial Contexts. Neuron *85*, 190–201. https://doi.org/10.1016/j.neuron.2014.12.001.
27. Maboudi, K., Giri, B., Miyawaki, H., Kemere, C., and Diba, K. (2024). Retuning of hippocampal representations during sleep. Nature *629*, 630–638. https://doi.org/10.1038/s41586-024-07397-x.
28. Ziv, Y., Burns, L.D., Cocker, E.D., Hamel, E.O., Ghosh, K.K., Kitch, L.J., El Gamal, A.E., and Schnitzer, M.J. (2013). Long-term dynamics of CA1 hippocampal place codes. Nat. Neurosci. *16*, 264–266. https://doi.org/10.1038/nn.3329.
29. Vaidya, S.P., Chitwood, R.A., and Magee, J.C. (2023). The formation of an expanding memory representation in the hippocampus. Preprint at bioRxiv. https://doi.org/10.1101/2023.02.01.526663.
30. Valero, M., Cid, E., Averkin, R.G., Aguilar, J., Sanchez-Aguilera, A., Viney, T.J., Gomez-Dominguez, D., Bellistri, E., and de la Prida, L.M. (2015). Determinants of different deep and superficial CA1 pyramidal cell dynamics during sharp-wave ripples. Nat. Neurosci. *18*, 1281–1290. https://doi.org/10.1038/nn.4074.
31. Cembrowski, M.S., Bachman, J.L., Wang, L., Sugino, K., Shields, B.C., and Spruston, N. (2016). Spatial Gene-Expression Gradients Underlie Prominent Heterogeneity of CA1 Pyramidal Neurons. Neuron *89*, 351–368. https://doi.org/10.1016/j.neuron.2015.12.013.
32. Dong, H.-W., Swanson, L.W., Chen, L., Fanselow, M.S., and Toga, A.W. (2009). Genomic–anatomic evidence for distinct functional domains in hippocampal field CA1. Proc. Natl. Acad. Sci. USA *106*, 11794–11799. https://doi.org/10.1073/pnas.0812608106.
33. Klausberger, T., and Somogyi, P. (2008). Neuronal Diversity and Temporal Dynamics: The Unity of Hippocampal Circuit Operations. Science *321*, 53–57. https://doi.org/10.1126/science.1149381.
34. Li, Y., Xu, J., Liu, Y., Zhu, J., Liu, N., Zeng, W., Huang, N., Rasch, M.J., Jiang, H., Gu, X., et al. (2017). A distinct entorhinal cortex to hippocampal CA1 direct circuit for olfactory associative learning. Nat. Neurosci. *20*, 559–570. https://doi.org/10.1038/nn.4517.
35. Somogyi, P., and Klausberger, T. (2005). Defined types of cortical interneurone structure space and spike timing in the hippocampus. J. Physiol. *562*, 9–26. https://doi.org/10.1113/jphysiol.2004.078915.
36. Hainmueller, T., and Bartos, M. (2018). Parallel emergence of stable and dynamic memory engrams in the hippocampus. Nature *558*, 292–296. https://doi.org/10.1038/s41586-018-0191-2.
37. Tononi, G., and Cirelli, C. (2003). Sleep and synaptic homeostasis: a hypothesis. Brain Res. Bull. *62*, 143–150. https://doi.org/10.1016/j.brainresbull.2003.09.004.
38. Wirtshafter, H.S., and Disterhoft, J.F. (2022). In Vivo Multi-Day Calcium Imaging of CA1 Hippocampus in Freely Moving Rats Reveals a High Preponderance of Place Cells with Consistent Place Fields. J. Neurosci. *42*, 4538–4554. https://doi.org/10.1523/JNEUROSCI.1750-21.2022.
39. Rolotti, S.V., Blockus, H., Sparks, F.T., Priestley, J.B., and Losonczy, A. (2022). Reorganization of CA1 dendritic dynamics by hippocampal sharp-wave ripples during learning. Neuron *110*, 977–991.e4. https://doi.org/10.1016/j.neuron.2021.12.017.
40. King, C., Henze, D.A., Leinekugel, X., and Buzsáki, G. (1999). Hebbian modification of a hippocampal population pattern in the rat. J. Physiol. *521*, 159–167.
41. Losonczy, A., Makara, J.K., and Magee, J.C. (2008). Compartmentalized dendritic plasticity and input feature storage in neurons. Nature *452*, 436–441. https://doi.org/10.1038/nature06725.
42. Mishra, R.K., Kim, S., Guzman, S.J., and Jonas, P. (2016). Symmetric spike timing-dependent plasticity at CA3–CA3 synapses optimizes storage and recall in autoassociative networks. Nat. Commun. *7*, 11552. https://doi.org/10.1038/ncomms11552.
43. Huerta, P.T., and Lisman, J.E. (1995). Bidirectional synaptic plasticity induced by a single burst during cholinergic theta oscillation in CA1 in vitro. Neuron *15*, 1053–1063. https://doi.org/10.1016/0896-6273(95)90094-2.
44. Hyman, J.M., Wyble, B.P., Goyal, V., Rossi, C.A., and Hasselmo, M.E. (2003). Stimulation in Hippocampal Region CA1 in Behaving Rats Yields Long-Term Potentiation when Delivered to the Peak of Theta and Long-Term Depression when Delivered to the Trough. J. Neurosci. *23*, 11725–11731. https://doi.org/10.1523/JNEUROSCI.23-37-11725.2003.
45. Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A., and Fiete, I. (2019). The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. Nat. Neurosci. *22*, 1512–1520. https://doi.org/10.1038/s41593-019-0460-x.
46. Park, S.P., Lopez-Rodriguez, F., Wilson, C.L., Maidment, N., Matsumoto, Y., and Engel, J., Jr. (1999). In vivo microdialysis measures of extracellular serotonin in the rat hippocampus during sleep-wakefulness. Brain Res. *833*, 291–296. https://doi.org/10.1016/s0006-8993(99)01511-5.
47. Nazari, M., Karimi Abadchi, J., Naghizadeh, M., Bermudez-Contreras, E.J., McNaughton, B.L., Tatsuno, M., and Mohajerani, M.H. (2023). Regional variation in cholinergic terminal activity determines the non-uniform occurrence of cortical slow waves during REM sleep in mice. Cell Rep. *42*, 112450. https://doi.org/10.1016/j.celrep.2023.112450.
48. Emrick, J.J., Gross, B.A., Riley, B.T., and Poe, G.R. (2016). Different Simultaneous Sleep States in the Hippocampus and Neocortex. Sleep *39*, 2201–2209. https://doi.org/10.5665/sleep.6326.
49. Durán, E., Oyanedel, C.N., Niethard, N., Inostroza, M., and Born, J. (2018). Sleep stage dynamics in neocortex and hippocampus. Sleep *41*, zsy060. https://doi.org/10.1093/sleep/zsy060.
50. Pedrosa, R., Nazari, M., Mohajerani, M.H., Knöpfel, T., Stella, F., and Battaglia, F.P. (2022). Hippocampal gamma and sharp wave/ripples mediate bidirectional interactions with cortical networks during sleep. Proc. Natl. Acad. Sci. USA *119*, e2204959119. https://doi.org/10.1073/pnas.2204959119.
51. Wierzynski, C.M., Lubenov, E.V., Gu, M., and Siapas, A.G. (2009). State-dependent spike-timing relationships between hippocampal and

prefrontal circuits during sleep. Neuron *61*, 587–596. https://doi.org/10.1016/j.neuron.2009.01.011.

52. Koike, B.D.V., Farias, K.S., Billwiller, F., Almeida-Filho, D., Libourel, P.-A., Tiran-Cappello, A., Parmentier, R., Blanco, W., Ribeiro, S., Luppi, P.-H., et al. (2017). Electrophysiological Evidence That the Retrosplenial Cortex Displays a Strong and Specific Activation Phased with Hippocampal Theta during Paradoxical (REM) Sleep. J. Neurosci. *37*, 8003–8013. https://doi.org/10.1523/JNEUROSCI.0026-17.2017.
53. Niethard, N., Brodt, S., and Born, J. (2021). Cell-Type-Specific Dynamics of Calcium Activity in Cortical Circuits over the Course of Slow-Wave Sleep and Rapid Eye Movement Sleep. J. Neurosci. *41*, 4212–4222. https://doi.org/10.1523/JNEUROSCI.1957-20.2021.
54. Bostock, E., Muller, R.U., and Kubie, J.L. (1991). Experience-dependent modifications of hippocampal place cell firing. Hippocampus *1*, 193–205. https://doi.org/10.1002/hipo.450010207.
55. Harris, K.D., Hirase, H., Leinekugel, X., Henze, D.A., and Buzsáki, G. (2001). Temporal interaction between single spikes and complex spike bursts in hippocampal pyramidal cells. Neuron *32*, 141–149. https://doi.org/10.1016/s0896-6273(01)00447-0.
56. Kesner, R.P., Farnsworth, G., and Kametani, H. (1991). Role of parietal cortex and hippocampus in representing spatial information. Cereb. Cortex *1*, 367–373. https://doi.org/10.1093/cercor/1.5.367.
57. Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999). Fast network oscillations in the hippocampal CA1 region of the behaving rat. J. Neurosci. *19*, RC20. https://doi.org/10.1523/JNEUROSCI.19-16-j0001.1999.
58. O'Neill, J., Senior, T., and Csicsvari, J. (2006). Place-Selective Firing of CA1 Pyramidal Cells during Sharp Wave/Ripple Network Patterns in Exploratory Behavior. Neuron *49*, 143–155. https://doi.org/10.1016/j.neuron.2005.10.037.
59. van der Maaten, L., and Hinton, G. (2008). Visualizing Data using t-SNE. J. Mach. Learn. Res. *9*, 2579–2605.
60. Mitra, P.P., and Pesaran, B. (1999). Analysis of dynamic brain imaging data. Biophys. J. *76*, 691–708. https://doi.org/10.1016/S0006-3495(99)77236-X.
61. Thomson, D.J. (1982). Spectrum estimation and harmonic analysis. Proc. IEEE *70*, 1055–1096. https://doi.org/10.1109/PROC.1982.12433.
62. Maboudi, K., Ackermann, E., de Jong, L.W., Pfeiffer, B.E., Foster, D., Diba, K., and Kemere, C. (2018). Uncovering temporal structure in hippocampal output patterns. eLife *7*, e34467. https://doi.org/10.7554/eLife.34467.
63. Dunn, B., Mørreaunet, M., and Roudi, Y. (2015). Correlations and Functional Connections in a Population of Grid Cells. PLoS Comput. Biol. *11*, e1004052. https://doi.org/10.1371/journal.pcbi.1004052.
64. Markus, E.J., Barnes, C.A., McNaughton, B.L., Gladden, V.L., and Skaggs, W.E. (1994). Spatial information content and reliability of hippocampal CA1 neurons: effects of visual input. Hippocampus *4*, 410–421. https://doi.org/10.1002/hipo.450040404.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental models: Organisms/strains		
Long Evans Rats	Janvier, France	RjOrl:LE
Software and algorithms		
KlustaKwik	Harris et al.[55]	https://github.com/klusta-team/klustakwik/
Python	Python Software Foundation	https://www.python.org
R	R Foundation	https://www.r-project.org/
Original code	This paper	https://doi.org/10.5281/zenodo.14803066
Other		
12um tungsten wires	California Fine Wire	M294520

EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

Three male Long-Evans rats (300-400 g, 2-4 months of age; Janvier, France) were used in the study. The animals were housed in a separate room on a 12-hour light/dark cycle and were taken to the recording room each day prior to the experiments. Animals shared a cage with littermates before surgery. All procedures involving experimental animals were carried out in accordance with Austrian animal law (Austrian federal law for experiments with live animals) under a project license approved by the Austrian Federal Science Ministry (License number: BMWFW-66.018/0015-WF/V3b/2014).

METHOD DETAILS

Surgery

Animals were implanted bilaterally with microdrives housing 32 (2x16) independently movable tetrodes targeting the dorsal CA1 region of the hippocampus. Each tetrode was fabricated out of four 10 μm tungsten wires (H-Formvar insulation with Butyral bond coat California Fine Wire Company, Grover Beach, CA) that were twisted and then heated to bind them into a single bundle. The tips of the tetrodes were then gold-plated to reduce the impedance to 200-400 kΩ. During surgery, the animal was under deep anesthesia using isoflurane (0.5%–3% MAC), oxygen (1-2l/min), and an initial injection of buprenorphine (0.1mg/kg). Two rectangular craniotomies were drilled centered above CA1 and positioned relative to bregma (centered at AP = -3.2; ML = ±1.6), the dura mater removed and the electrode bundles implanted into the superficial layers of the neocortex, after which both the exposed cortex and the electrode shanks were sealed with paraffin wax. Five to six anchoring screws were fixed onto the skull and two ground screws (M1.4) were positioned above the cerebellum. After removal of the *dura*, the tetrodes were initially implanted at a depth of 1-1.5 mm relative to the brain surface. Finally, the microdrive was anchored to the skull and screws with dental cement (Refobacin *Bone Cement R*, Biomet, IN, USA). Two hours before the end of the surgery, the animal was given the analgesic Metacam (5mg/kg). After a one-week recovery period, tetrodes were gradually moved into the dorsal CA1 cell layer (*stratum pyramidale*). After completion of the experiments, the rats were deeply anesthetized and perfused through the heart with 0.9% saline solution, followed by a 4% buffered formalin phosphate solution for the histological verification of the electrode tracks.

Data Acquisition, Training and Behavior

The animals were housed individually in a separate room under a 12h light/12h dark cycle. Following the postoperative recovery period, rats were reduced to and maintained at 85% of their age-matched preoperative weight. Water was available *ad libitum*. Each animal was handled and familiarized with the recording room and with the general procedures of data acquisition. Behavioral training was performed after electrode implantation during days when the electrodes were moved towards the hippocampus but before they reached the hippocampus. Overall recordings were performed in three rats that were trained to perform the seek of the hidden rewards task on the cheeseboard maze[9,56] and come back to the start-box. In order to achieve this, random groups of visible food pellets (MLab rodent pellet 20 mg, TestDiet) were spread out on the surface of the cheeseboard maze while the rat was inside the start box. Then we opened the door and left the animal freely foraging the entire maze and once the animal returned to the start box, we closed the start-box door. With the help of this training protocol, we could shape our animals' behavior to explore the entire maze and return to the start box automatically. Despite the automatic behavior, we could ensure that, under experiments, rats had no or limited experience in performing the cheese-board maze task during the time of the recordings. Each daily experiment consisted of a sequence of nine recording sessions in the following order: a free exploration session on a familiar environment, half-hour immobility/sleep rest session in the animal's own cage, free exploration on the cheese-board, an immobility/sleep rest session (own cage), a learning session (4 randomly selected invisible locations) on the cheese-board, 20-hour continuous monitoring in the cage, recall learning session with the same bait locations, an immobility/sleep rest session (own cage), post probe on the cheese-board (free exploration, unrewarded) and another free exploration session on the same familiar environment as the first one. During the learning session, once animals learned the four invisible goal locations, they performed only five additional trials to master the task. Following the ~20-hour long rest session, we tested the recall performance of the animal on the cheeseboard maze in 50 further rewarded trials with the same bait locations, which were learned 20 hours before.

To be able to record extracellular electric signals continuously over a long period of time, we adapted to our commonly built 3D printed microdrive a new, high-fidelity 64- (two animals) and 128-channel (one animal) wireless recording system from TBSI (Triangle BioSystems, Durham, NC, W128) to use in our experiments. In animals in which only the 64-channel telemetry system was available, only half of the recorded channels of the 128-channel drive were used. Using this experimental preparation, we recorded cell population activity continuously over 30 hours during learning, long periods of rest where reactivation takes place and memory recall tests at the end. This telemetry system has been able to amplify and transmit wide band (0.7 Hz to 9 kHz) signals above 20 kHz on the analog channels, which were then digitized at 20 kHz.

Two small light-emitting diodes (LEDs) mounted on the preamplifier headstage were used to track the location of the animal via an overhead video camera. The animal's location was constantly monitored throughout the experiment. We detected the two LEDs with custom-made tracking software (*positrack*, github.com/kevin-allen/positrack) made by Kevin Allen. The video signal has been triggered and tracked continuously with a TTL pulse sent by the camera's computer on a common analog channel.

QUANTIFICATION AND STATISTICAL ANALYSIS

Spike Detection, Sorting and Stability

The spike detection and sorting procedures, and clustering were performed as previously described.[57,58] Continuously recorded wide-band signals were digitally high-pass filtered (0,8-5 kHz). Action potentials were extracted by first computing power in the 800-9000 Hz range within a sliding window (12.8 ms). Action potentials with a power of > 5 SD from the baseline mean were selected and spike features were then extracted by using *principal components analyses* (PCA). The detected action potentials were segregated into multiple single units by using automatic clustering software (http://klustakwik.sourceforge.net/). These clusters were manually refined by a graphical cluster-cutting program. Only units with clear refractory periods in their autocorrelation (<2 ms) and well-defined cluster boundaries were used for further analysis. We further confirmed the quality of cluster separation by calculating the Mahalanobis distance between each pair of clusters.[55] To be able to analyze changes in the firing patterns of neuronal ensembles over time, we had to guarantee that our set of putative cells was sampled from clusters with stable firing over the whole recording. To ensure this, we clustered together periods of waking and rest sessions and then we plotted spike features over time by plotting 2-dimensional unit PCA cluster plots across the whole recording in addition to the stability of spike waveforms. With the help of this method, we could exclude those spike clusters that overlapped during the course of recordings. To further verify spike cluster stability, we used the *t-student stochastic neighbor embedding* (t-SNE) dimensionality reduction method: t-SNE embeds the n-dimensional extracellular spikes (n = number of features by which each spike is decomposed) into a low dimensional space.[59] The t-SNE focuses on ensuring that the local structure remains intact while it ignores the global structure; therefore, when we expressed t-SNE features over time, we could visually exclude those clusters that were unstable during the whole recording due to electrode drifting.

Putative pyramidal cells and putative interneurons in the CA1 region were discriminated by their autocorrelations, firing rate, and waveforms, as previously described.[57]

Sleep classification

In recordings, exploratory and immobility or quiet rest periods were manually separated offline as previously described.[57,58] For each session, the theta/delta ratio was plotted against speed so that the behavioral state could be manually identified. The theta/delta field power ratio was measured in 1,600 ms segments (800 ms steps between measurement windows) with Thomson's multitaper method.[60,61] Waking behavior included periods of locomotion and/or the presence of theta oscillations (visible in the theta/delta ratio), with no more than ~2.5 s of transient immobility. Rest epochs were selected when both the speed and theta-delta ratio dropped below a pre-set threshold (speed: <4cm/s, theta/delta ratio: <2) for at least 2.4 s. During periods of active waking behavior, theta-oscillatory wave detection was performed as previously described[57,58] using the negative peaks of individual theta waves from the filtered trace of the local field potential (5–28 Hz). The band used for the detection was wider than the theta band in order to precisely detect the negative peaks of the theta waves, which otherwise would have smoothed out by using a narrow theta band. Quiet rest segments have been identified by longer periods of immobility (because of the lengths of our recording, at least 10 min) and the clear presence of REM-theta and slow-wave field oscillations.

Sharp wave ripple detection

For the detection of SWRs, local field potentials were band-pass filtered (150–250 Hz), and a reference signal (to ensure the lack of ripple activity, we left a tetrode above the hippocampus as a reference) was subtracted to eliminate the so-called common noise (muscle artifacts due to scratching, twitching, etc.). The power (root mean square) of the filtered signal was calculated for each electrode and summed across electrodes considered to be in the CA1 *stratum pyramidale*. The threshold for SWR detection was set to 7 SD above the background mean. The SWR detection threshold was always set in the first immobility session, but longer (at least 35 min) as it was described earlier and the same threshold was used for all other sessions.[57,58]

Quantification of stability of clustering features over time

To assess the temporal stability of waveforms, we used the first 12 clustering features. For each feature, the mean was computed for the corresponding interval (5th, 10th, 15th hour) and z-scored using the mean and std of the first hour.

Cell separation over time

We estimated the separation of cells by using the first 12 clustering features and three different approaches. First, for each cell, we divided the distance (1 - Pearson R) between the mean value of each feature during the nth hour and the cell's own mean feature values from the first hour by the distance between the mean feature values during the nth hour and the other cells' feature means from the first hour. Using this measure we investigated whether a cell's waveform at a later stage during the experiment is more similar to its own waveform from the first hour than to first hour waveforms from other cells. Second, for each cell, we divided the distance (1 - Pearson R) between the mean feature values during the nth hour and the cell's own mean values from the first hour by the distance between the other cells' feature means during the nth hour and the cell's own feature means from the first hour. This measure estimates whether other cells' waveforms at later stages during the experiment are more similar to the cell's waveform from the first hour than its own waveforms at later stages. Third, we compared the across-cell distance with the within-cell distance. For the across-cell distance, the distance between the mean feature values per cell of the first hour and the other cells' feature means from the first hour was calculated. The within distance corresponds to the distance between the cell's feature averages from the first hour and the same cell's feature averages of the last hour of the experiment.

Clustering feature stability

For each cell, we calculated the mean and std for each of the first 10 PCA clustering features using data from the first hour. We then computed the mean per clustering feature of each cell at different time intervals (7th, 15th, and 21st hour) of the experiment. In order to test whether clustering features drift away from the initial values, we z-scored the means during different time intervals using the mean and std from the first hour.

Measuring excess path during acquisition and recall

We assessed the animal's ability to learn and recall goal locations on the cheeseboard by computing the excess path: once the rat had left the start box, we measured the length of the path the animal took to reach any of the four goals (animal position within 10 cm radius around goal location). Next, we detected when the animal left the goal again and measured the path length to the next goal. This procedure was repeated for the remaining two goals. We then calculated the optimal paths as straight lines between either the start location and the first goal or between subsequently visited goals. Each taken path length was then divided by the optimal path length to yield the excess path as a multiple of the optimal path.

Poisson Hidden Markov Model (pHMM) and model fitting

We trained two separate hidden Markov models with Poisson emissions (pHMM) on the neural data obtained during the cheeseboard task before (acquisition) and after (recall) the long sleep. Only data from running periods (speed > 5cm/s) was used. The acquisition data length was matched in terms of duration to the recall data to have the same training data length for acquisition and recall. Then, the neural data was binned using temporal bins of 100 ms length.

The pHMM model assumes the temporal evolution of an unobserved discrete state as described in Maboudi et al.[62] In short, the probability of observing an ensemble E_N of N independently firing neurons at time t for state i can be modeled as:

$$P(state = i) = \prod_{n=1}^{N} P(state = i) \propto \prod_{n=1}^{N} (\lambda_{n,i})^{s_{nt}} \, exp(-\lambda_{n,i}) \qquad \text{(Equation 1)}$$

where s_{nt} is the firing rate of neuron n at time t. The firing rate is modeled according to a Poisson process with a mean firing rate $\lambda_{n,i}$ defined by the unobserved discrete state i.

The transition probabilities between M unobserved states were captured by the M x M transition matrix **A**. The hyperparameter M defining the number of states of the model was determined using the cross-validated maximum likelihood (Figure S2). All model parameters were computed using the EM algorithm.

Rate map generation for Bayesian decoding

For each session and cell, one rate map for acquisition and one rate map for recall were computed. In order to reconstruct the two-dimensional spatial distribution of each cell's firing probability, we used a maximum entropy model inference paradigm.[63] Only running periods were considered by applying a 5 cm/s speed filter. The cell activity was binned using time windows of 10 ms.

We assigned a binary variable $S_i(t)$ to each neuron to denote the presence or absence (+1/-1) of spikes emitted by neuron *i* within time bin *t*. Including the dependence of each neuron's state on the previous state (*t*-1) of the population, the maximum entropy distribution over the state $S_i(t)$ of neuron *i* at time *t* is

$$p(S_i(t)) = \frac{exp[S_i(t)H(t-1)]}{2\cosh[H(t-1)]} \qquad \text{(Equation 2)}$$

with H(*t*) being the time-varying covariate representing the external field in statistical physics. Equation 1 describes a Generalized Linear Model (GLM) with the interaction kernel extending to a one-time step in the past.

By maximizing the log-likelihood with respect to H(*t*) we found the most likely values of H(*t*) to generate the observed data:

$$L[S,H] = \sum_i [S_i(t)H_i(t-1) - log(2\cosh H_i(t-1))] \qquad \text{(Equation 3)}$$

For the spatial input we assumed the sum of two-dimensional Gaussian basis functions centered on an evenly spaced M x M square lattice that spanned the cheeseboard. The spatial field of cell *i* at time *t* can then be computed in the following way:

$$H_i(t) = \sum_{j=1,k=1}^{M,N} a_{ijk}\, exp\left[\frac{-\left((x(t)-x_{jk})^2+(y(t)-y_{jk})^2\right)}{r^2}\right]+h_i \qquad \text{(Equation 4)}$$

where h_i is the baseline activity of cell *i*, (x_{jk}, y_{jk}) are the centers of the Gaussians and *r* is the standard deviation of basis functions. An accurate representation of the cell activity in space was computed by inferring parameters α_{ijk} of the linear combinations of Gaussian basis functions. The resulting map was then partitioned into bins of 4 cm.

Decoding rest activity using pHMM

In order to compensate for differences in temporal dynamics between REM and NREM rest, we binned the rest data using bins with a constant number of 12 spikes. Since the awake pHMM models were trained on temporal bins of 100 ms, we computed a scaling factor between awake and rest neural activity to match the two. First, we calculated the mean number of spikes occurring within 100 ms time bins during awake behavior $n_{\text{awake},100\text{ms}}$. The scaling factor γ_{phmm} is defined as:

$$\gamma_{phmm} = \frac{12}{n_{awake,100ms}} \qquad \text{(Equation 5)}$$

The likelihood of the rest activity at time *t* given the discrete pHMM state *i* is computed as follows:

$$L(state = i) = \prod_{n=1}^{N}\frac{\gamma_{phmm}\lambda_{i,n}^{s_{n,t}}}{s_{n,t}}exp(\gamma_{phmm}\lambda_{i,n}) \qquad \text{(Equation 6)}$$

where $s_{n,t}$ is the number of spikes of neuron *n* at time *t* and $\lambda_{i,n}$ is the mean firing rate of that neuron in state *i*. Notice that for our decoding procedure, the transition probabilities across states were not considered. To assess which state was most likely reactivated at time *t* during rest, the state with the maximum log-likelihood was selected.

Decoding rest activity using Bayesian decoding

Equivalent to the pHMM decoding approach, the rest data was binned using bins with a constant number of 12 spikes. In the case of Bayesian decoding, the acquisition and recall rate maps were computed using 10 ms time bins. Therefore, the scaling factor γ_{bayesian} is computed using the mean number of spikes occurring within 10 ms time bins $n_{\text{awake},10\text{ms}}$ during awake behavior:

$$\gamma_{bayesian} = \frac{12}{n_{awake,10ms}} \qquad \text{(Equation 7)}$$

The likelihood of the rest activity at time *t* given the spatial bin *x* is given by:

$$L(spatial\ bin = x) = \prod_{n=1}^{N}\frac{\gamma_{bayesian}\lambda_{x,n}^{s_{n,t}}}{s_{n,t}}exp(\gamma_{bayesian}\lambda_{x,n}) \qquad \text{(Equation 8)}$$

where $s_{n,t}$ is the number of spikes of neuron *n* at time *t* and $\lambda_{x,n}$ is the mean firing rate of that neuron in spatial bin *x*. To reconstruct the spatial bin that was most likely reactivated at time *t*, the spatial bin with the maximum log-likelihood was selected.

Quantifying drift during rest

For each time point *t* in rest, we calculated the Drift score in the following way:

$$Drft\ score = \frac{max(L_{recall}) - max(L_{acquisition})}{max(L_{recall})+max(L_{acquisition})} \qquad \text{(Equation 9)}$$

with L_{recall} and $L_{acquisition}$ being the maximum likelihoods across all states or spatial bins (when the Bayesian decoding was used in Figures S3H–S3J) for the acquisition or recall models (pHMM or rate maps) at time t, respectively. The resulting Drift score was smoothed across time.

First vs. second half of rest

In order to validate whether the drift is more prominent during the first half of the rest, we computed the delta of the drift score for the first and second half of the rest separately. Then, we tested whether the ratio between the delta of the first half and the delta of the second half was greater than 1 using the student T-test.

Net effect vs. cumulative effect of drift

In order to assess the amount of memory drift with respect to different timescales, we compared the net effect and cumulative effect. For the net effect, we computed the difference in drift score between the beginning and the end of the rest period. On the other side, the cumulative effect was computed by summing the absolute values of the drift score throughout the rest period.

Opposing effects on drift in NREM and REM

NREM and REM periods were identified as described above. Neighboring rest epochs of the same type were merged to obtain a set of alternating NREM/REM epochs. For each epoch, the change in drift score was computed by subtracting the first value of the epoch from the last value of the epoch.

Calculation of firing rate, LFP amplitude, and SWR number per REM/NREM epoch

To relate the firing rate and LFP amplitude to delta drift score, we measured the firing rate (z-scored across the entire session) and average LFP amplitude per epoch and we also noted the delta drift score for that specific epoch and correlations were calculated between delta drift score vs firing rate and delta drift score vs LFP amplitude during REM and NREM epochs. Similarly, SWR numbers were calculated for each NREM epoch and correlated with the delta drift score.

Persistent and unstable subsets

Persistent and unstable cell subsets were identified based on changes in their firing rate distributions from acquisition to recall. For each cell, the distributions of firing rates for acquisition and recall were computed separately. If the acquisition distribution was significantly greater than the recall distribution (one-sided MWU, $p<0.01$), the cell was assigned to the decreasing subset. If, on the other hand, the recall distribution was significantly greater than the acquisition distribution (one-sided MWU, $p<0.01$), the cell was labeled as increasing. All other cells that did not show a significant difference in their firing rate distributions from acquisition to recall made up the persistent subset.

Drift using different subsets of cells

To evaluate the effect of using only a subset of cells for our rest decoding procedure on the observed drift we proceeded as follows. We removed all cells not contained in the subset from state vectors of the original pHMM model and computed the drift score using the maximum likelihoods from the acquisition and recall states. Then, the drift score for the entire rest duration was to fit a line and the resulting slope was calculated.

Firing probability changes in NREM and REM

REM and NREM epochs were identified as described above. Thereafter, for each subset of persistent, increasing and decreasing cells, we computed the change in firing probability from the beginning to the end of each epoch. The firing probability is defined as the number of spikes contributed by the subset to the total number of spikes per constant spike bin. We computed the change by subtracting the firing probability of the first bin of the epoch from the value of the last bin of the epoch. Only epochs with significant changes in firing probability were considered.

Measuring spatial information

To assess the spatial information of single cells we computed the sparsity and spatial information per second as previously described.[64] The spatial information per second was computed using the following equation:

$$I_{sec} = \sum_{i=1}^{M} p_i \lambda_i \log\left(\frac{\lambda_i}{\lambda}\right) \quad \text{(Equation 10)}$$

where p_i and λ_i are the probability of occupying and the firing rate of bin i, respectively. Parameter λ describes the mean firing rate of the cell in the environment.

Decoding positions using neural activity during behavior

Bayesian decoding. We applied standard Bayesian decoding using 1 cm^2 spatial bins. First, we computed the mean number of spikes $\lambda_{n,i}$ for each cell n and spatial bin i representing all the bins of the two-dimensional environment. Given the spikes S_N of N neurons at time t, we computed the likelihood of being in bin i using the following equation:

$$L(bin = i) = \prod_{n=1}^{N} \frac{\lambda_{n,i}^{s_n} \exp(-\lambda_{n,i})}{s_n!} \quad \text{(Equation 11)}$$

with s_n being the number of spikes of neuron n at time t. The spatial bin with the highest likelihood represented the decoded location.

Decoding position using pHMM. Using the entire neural data from the acquisition and the trained pHMM, we inferred the most likely state sequence using the Viterbi algorithm. By matching the sequence of states with the tracking data of the animal, we identified a mean spatial location for each state.

Given the activity of N neurons at time t, we computed the normalized likelihood for each state of our pHMM. The decoded location was then calculated by weighing the mean location of each state with its normalized likelihood and computing the average position across all states.

Computing mean firing rates

Acquisition, rest and recall were split into 5 min. chunks to compute the mean and maximum firing rates of the different cell subsets.

Distance between peak firing and closest goal

For each cell, we computed its rate map and determined the location on the maze with the highest firing rate. Next, we calculated the distance between the location with the peak firing rate and the closest goal location.

Population vector correlations

Neurons were separated into persistent, increasing and decreasing cells according to the procedure defined above. Neural activity per subset was binned using spatial bins of 10 cm^2 size, yielding one population vector per spatial bin. Spatial bins that were not visited in all relevant behavioral episodes were excluded. Then, Pearson correlations between population vectors of the same spatial bin during the different behavioral episodes were computed.