

Replay of rule-learning related neural patterns in the prefrontal cortex during sleep

Adrien Peyrache¹, Mehdi Khamassi^{1,2}, Karim Benchenane¹, Sidney I Wiener¹ & Francesco P Battaglia^{1,3}

Slow-wave sleep (SWS) is important for memory consolidation. During sleep, neural patterns reflecting previously acquired information are replayed. One possible reason for this is that such replay exchanges information between hippocampus and neocortex, supporting consolidation. We recorded neuron ensembles in the rat medial prefrontal cortex (mPFC) to study memory trace reactivation during SWS following learning and execution of cross-modal strategy shifts. In general, reactivation of learning-related patterns occurred in distinct, highly synchronized transient bouts, mostly simultaneous with hippocampal sharp wave/ripple complexes (SPWRs), when hippocampal ensemble reactivation and cortico-hippocampal interaction is enhanced. During sleep following learning of a new rule, mPFC neural patterns that appeared during response selection replayed prominently, coincident with hippocampal SPWRs. This was learning dependent, as the patterns appeared only after rule acquisition. Therefore, learning, or the resulting reliable reward, influenced which patterns were most strongly encoded and successively reactivated in the hippocampal/prefrontal network.

The acquisition of labile new memories can trigger processes spanning from molecular¹ to system-wide levels, gradually transforming and stabilizing memory traces. The system consolidation theory views the interaction between hippocampus and neocortex as being instrumental for this^{2–4}. Although the hippocampus is vital in the initial acquisition and early storage of memories, the cerebral cortex, among other structures, is important later on⁵. The exchange between a fast-learning module (the hippocampus) and a slower one (the neocortex) would take place mainly after memory acquisition, allowing one-shot acquisition of new items without losses of older memories because of interference^{2,3}. A further role of slow consolidation following acquisition would be to reorganize memories into more semanticized, de-contextualized representations^{6–8}.

A role for SWS in such an exchange^{3,5,9–11} would be to replay the neural patterns concerning previously acquired information^{12–18}. Such sleep replay would then instill a change in the neural substrate of memory traces and ultimately favor memory consolidation. During sleep, the hippocampus and the neocortex engage in a dialog that involves and affects the dynamical states of both^{19–23}. Hippocampal SPWRs are likely vectors for hippocampal-neocortical information exchange²⁴; SPWRs²⁵ are brief (~50–150 ms), large bursts of hippocampal activity that are mostly observed during SWS or immobility and correspond to increased hippocampal memory reactivation¹⁴. During SWS, neocortical activity shows periods of large, synchronous oscillations (0.1–4 Hz) of membrane potentials and neural firing²⁶, and these are correlated with SPWRs^{19–21}. Slow oscillations were recently found to coordinate episodes of visual cortical and hippocampal

reactivation¹⁶, but the precise temporal relationship between cortical and hippocampal replay remains unknown.

The PFC is often implicated in long-term memory consolidation²⁷, particularly for hippocampally dependent spatial and contextual information. Indeed, the PFC shows detailed, time-compressed replay following initial acquisition of memory-related sequences of neural ensemble activation in rats¹⁷ and increased coordination with the hippocampus during retrieval of sleep-consolidated memories in humans²⁸. The PFC is one of the neocortical areas most closely associated with the hippocampus, both anatomically and physiologically, as it has a unique afferent pathway from the hippocampus²⁹ that is endowed with synaptic plasticity³⁰. Some functional imaging and immediate early gene expression data support the idea that the hippocampus activity contributions decrease over time during consolidation, with an opposite, increasing trend being observed for the PFC^{27,28,31}. However, the concerted function of PFC and hippocampus is also necessary for memory maintenance during task performance^{32,33}.

Although the behavioral electrophysiology literature provides numerous examples of memory replay^{12,14–18}, the animals in these studies were over-trained and little learning actually took place, or no specific analysis of the evolution of replay with task performance was attempted¹³. Our goal was to investigate memory reactivation and hippocampal-neocortical interactions while new task-relevant information was actually being acquired to better characterize the link between learning and memory replay processes. Moreover, we focused on subsecond resolution of the time course of memory replay to

¹Laboratoire de Physiologie de la Perception et de l'Action, Collège de France, Centre National de la Recherche Scientifique, Paris, France. ²Institut des Systèmes Intelligents et de Robotique, Université Pierre et Marie Curie – Paris 6, Centre National de la Recherche Scientifique, Paris, France. ³Center for Neuroscience, Swammerdam Institute for Life Sciences, Faculty of Science, Universiteit van Amsterdam, Amsterdam, The Netherlands. Correspondence should be addressed to S.I.W. (sidney.wiener@college-de-france.fr).

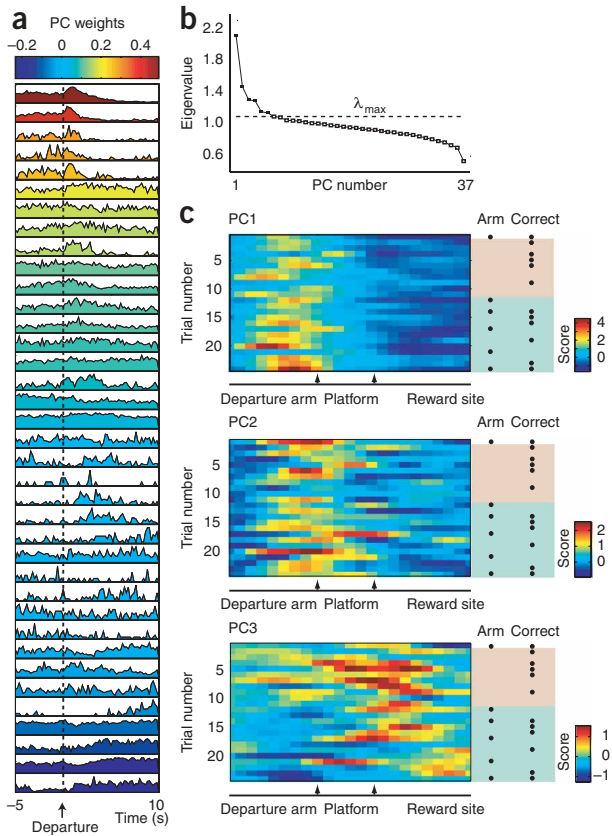


Figure 1 Signal components and their behavioral correlates. **(a)** Peri-event time histograms of all cells from a recording session, aligned with trial initiation (vertical dashed line). Cells are sorted by their weight in the first signal component (scale above); cells with similar weights tended to have similar behavioral correlates and cells with oppositely signed large weights had complementary behavioral correlates. PC, principal component.

(b) Awake epoch correlation matrix eigenvalues from the same session as is shown in **Figures 2a** and **3a**. The dotted line represents the signal threshold (λ_{max}), defined as the theoretical upper bound for eigenvalues in the case of random spike trains (see Online Methods). Filled squares indicate eigenvalues associated with the six (suprathreshold) signal components. Hollow squares represent nonsignal components. **(c)** Trial-by-trial time scores of the first three principal components during the awake epoch, plotted as a function of the linearized position of the rat on the maze.

The right panels summarize the rat behavior. The beige background indicates trials in which the rat reliably chose the right arm, and the blue background indicates trials in which the rat basically alternated between the two arms (with no other discriminable strategy). The black dots in the first column (Arm) denote trials in which the rat chose the left arm. The black dots in the second column (Correct) indicate rewarded trials (here the lit arm was rewarded and neither strategy was successful). Note that principal component scores are shown instead of reactivation strength (see below) because in this case the sign is important, as large positive principal component scores denote activation of those cells with positive principal component weights (see **Fig. 2a**) and large negative principal component scores denote activation of those cells with negative principal component weights.

study precise correlations between replay events and large-scale synchronization phenomena during SWS, such as SPWRs and slow oscillations. Learning-related changes in neural activity over brief time scales have been described in both the prefrontal cortex^{34,35} and hippocampus, but the effects of these changes on subsequent sleep activity has not yet been studied. We recorded neural activity in the PFC and the hippocampus in rats during a cross-modal rule shift task (known to implicate the medial PFC³⁶), which allowed us to introduce novel elements in the form of new rules while leaving the perceptual aspects of the task unchanged.

RESULTS

mPFC ensemble activity patterns during a rule shift task

We used multiple tetrodes to record ensembles of medial PFC (see Online Methods and **Supplementary Fig. 1** online) neurons together with mPFC and hippocampal local field potentials (LFPs) in rats. The animals performed a task in a Y maze (**Supplementary Fig. 2** online), in which they had to learn to select the rewarded arm using one of four possible rules (left arm, right arm, illuminated arm and non-illuminated arm; during each trial, one target arm was illuminated at random). This period will be referred to as the awake epoch. Neural activity was also monitored during rest periods immediately before and after the awake epoch (pre and post epochs). As soon as the rat achieved criterion performance (see Methods) according to the current rule, the rule was changed without any additional cue and the rat had to again infer the new rule from the pattern of rewarded and nonrewarded arms. Because no pre-training was performed before the electrophysiological recordings, during the experiments the rats encountered rules to which they had never been exposed before.

We recorded 1,692 cells in the mPFC (**Supplementary Fig. 1**) from four rats during a total of 63 recording sessions (rat 15, 16; rat

18, 11; rat 19, 12; rat 20, 24). Only sessions with a minimum of ten cells and at least 4 min of SWS in each rest epoch were analyzed. Cells in the mPFC had diverse behavioral correlates, corresponding to one or more task phases, and responses in some neurons dynamically adapted as the rat acquired the current task rule (Battaglia *et al.*, *Soc. Neurosci. Abstr.* **573.13**, 2006). We used principal component analysis to extract the neural patterns characteristic of the awake epoch (high-rank principal components, associated with larger eigenvalues or encoding strengths will be referred to as signal components, whereas lower-rank, nonsignal components mostly reflect noise; see Online Methods and **Supplementary Fig. 3** online).

Signal components identified neuronal assemblies with reliable and consistent responses in the task. For example, they assigned same-signed weights to cells with similar behavioral correlates and opposite-signed weights to cells with complementary correlates (**Fig. 1a** and **Supplementary Fig. 4** online). On the basis of the eigenvalues associated with the principal components and on a threshold value computed from the null hypothesis of random, uncorrelated spike trains, we could typically discriminate 1–6 signal components (and occasionally more) in each session (**Fig. 1b**). The patterns of activity detected by principal components were correlated with behavior; in an example session, the first principal component showed a positive peak activation right after trial onset (**Fig. 1c**), principal component 2 peaked later in the trial and principal component 3 peaked even later on. Moreover, principal components 1 and 2 increased their scores, whereas principal component 3 decreased its score as the rat abandoned the strategy of always going to the right arm across trials (**Fig. 1c**) and instead chose, with a great probability, to alternate between the two target arms. Thus, principal components 1, 2 and 3 extracted patterns of activity that correlated both with trial phase and, at a greater time scale, with the strategy that the rat followed in a block of trials.

Transient synchronized replay of awake patterns

To assess the nature and extent of the interaction between prefrontal cortex and hippocampus in memory replay, we characterized the

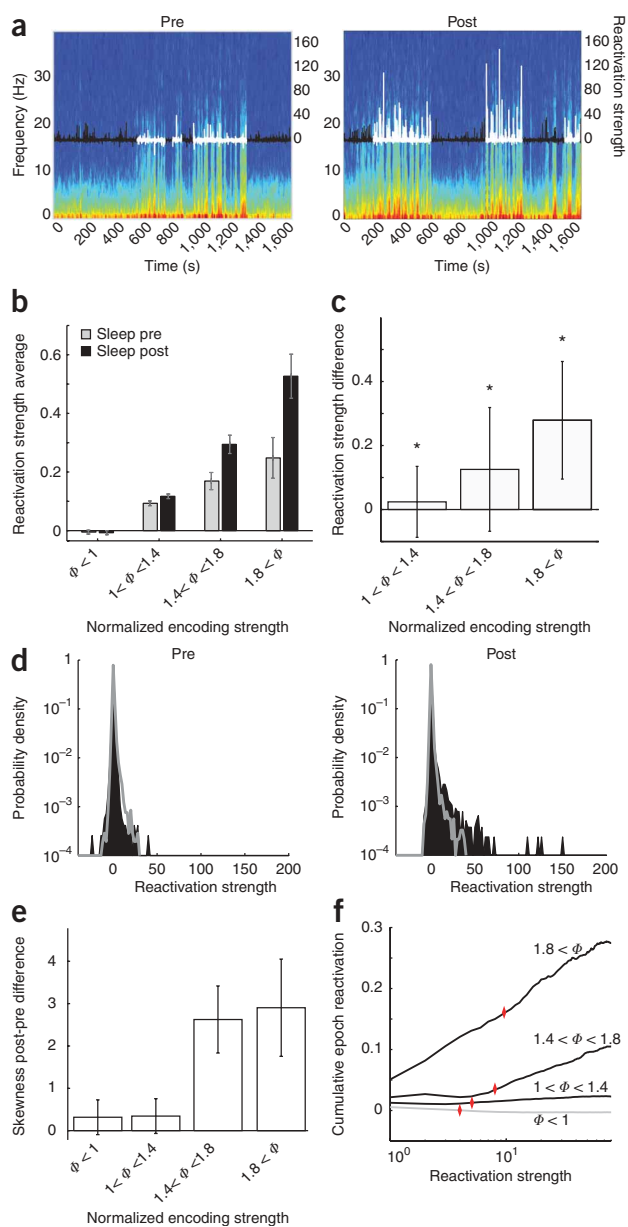


Figure 2 Time course of memory replay. **(a)** Reactivation strength (black and white traces; right axes) of the signal component during sleep pre (left) and post (right) epochs superimposed on the mPFC LFP spectrogram (left axes). White traces indicate SWS periods and black traces indicate non-SWS. The spectrogram shows periods of elevated slow/delta and spindle oscillations often coinciding with SWS. Reactivation strengths show high peaks during post epoch SWS, usually concomitant with periods of strong oscillations. **(b)** Bar plot of the epoch-wide average reactivation strengths, grouped by their encoding strength (normalized eigenvalue), $\Phi = \lambda/\lambda_{\max}$, for the pre and post epochs (data from 63 sessions and 1,692 recorded cells). Error bars indicate s.e.m. **(c)** Average difference in reactivation strength between pre and post epochs for principal components grouped by their encoding strengths. Error bars indicate s.d. **(d)** Incidence of reactivation strengths during rest periods (left, pre; right, post) for the same session as in **a** (also shown in **Fig. 3a**). Black filled zones indicate SWS periods and the gray trace indicates non-SWS. The post epoch SWS histogram has a heavy tail, reflecting strong transient reactivation events. **(e)** Average difference between the skewness of post and pre epoch reactivation strength incidence histogram (**Fig. 3c**) for components grouped by their encoding strength; post epoch histograms were generally more skewed than pre epoch ones. Error bars represent s.e.m. **(f)** Cumulative contribution to the difference between total reactivation strength in post and pre epochs for signal components (black, grouped by encoding strength) and for nonsignal components (gray). Red diamonds indicate the 99th percentile of reactivation strength distribution (in post). About half of the reactivation was accounted for by reactivation strengths over the 99th percentile.

the three signal groups observed here, sorted according to their encoding strength, $n = 811$ for nonsignal components; **Fig. 2b,c**). Reactivation strength correlated positively with encoding strength ($r^2 = 0.61$, $P < 10^{-30}$, Pearson correlation test, $n = 323$; **Fig. 2b** and **Supplementary Fig. 7** online). Thus, the patterns that were most active in the mPFC during the awake epoch were preferentially reactivated during the following SWS, similarly to previous observations in the hippocampus¹⁸. During pre SWS, this relationship was significantly weaker with respect to slope ($P < 10^{-20}$, $n = 323$) and correlation ($P < 10^{-5}$, $n = 323$; **Supplementary Fig. 7**). These observations were not likely to erroneously result from potentially faulty spike sorting, as they persisted when cell pairs discriminated from the same tetrode were ignored. Moreover, cell pairs from the same tetrodes, when considered alone, showed no replay effect, so that virtually all contributions to the replay results came from the correlations between neurons recorded with different tetrodes (**Supplementary Fig. 8** online).

Notably, replay occurred in distinct events of strong signal reactivation in post epoch SWS (**Fig. 2a**), denoting synchronous transient activation of the cell assemblies identified by the signal components. Histograms of the reactivation strengths for post epoch SWS were heavy tailed (**Fig. 2d**), with the tail constituting the main difference with pre epoch SWS (**Fig. 2d**), whereas the bulk of the distribution was similar in the pre and post epochs. The difference in tail weight was reflected in the significant difference between the post and pre epochs in the skewness of the signal reactivation–strength histograms ($P < 0.05$, t test, $n = 10$ – 40), most markedly for patterns with higher encoding strengths (**Fig. 2e**). The peaks in reactivation strength correspond to the transient, coordinated activation of the cells that are assigned a large weight in the relative principal component (**Supplementary Fig. 9** online). Those cells provide the greatest contribution to the total reactivation strength. Different principal components recruit different, rarely overlapping sets of neurons with high-weight strengths. During sleep, reactivation strengths for simultaneously recorded principal components tended not to peak at the same times; instead, concomitant activation of different principal component–related patterns was less than expected by chance, as can be inferred from the zero lag trough in their cross-correlograms

detailed time course of replay during rest episodes. For this, we computed the instantaneous reactivation strength (see Methods) of the signal components computed from the awake epoch. At each moment (with a resolution of 100 ms, unless otherwise specified), reactivation strength assesses the similarity between reference awake signal components and the rest period neural activity.

During post epoch SWS, signal components reappeared more frequently and strongly than in the pre epoch (for example, **Fig. 2a**), confirming that experience-related patterns are reactivated in mPFC¹⁷ in ensuing sleep. No such effects were observed in the rest periods that were not classified as SWS (as shown in **Supplementary Fig. 5** online); therefore, further analyses were restricted to SWS. Pre and post epochs SWS did not differ in terms of the average duration of the sleep episodes, average population firing rates, rates of occurrence of delta waves and SPWRs, and local field potential power in the delta and spindle ranges (**Supplementary Fig. 6** online). The average reactivation strength was greater during post epoch SWS than pre epoch SWS for signal components ($P < 0.005$ all comparisons, $n = 10$, 40 and 273 for




Figure 3 Memory replay, SPWRs and cell activity in a rest session. (a) Reactivation strength for one signal component. (b) Occurrences of SPWRs. (c) Averaged firing rate of the cells that contributed the highest weights to the signal component. Note that peaks correspond to events in a and b. (d) Average firing rate for all cells recorded in this session.

(Supplementary Figs. 9 and 10 online). Shuffled controls showed that this cannot be explained by global fluctuations in the population firing rate alone (Supplementary Fig. 11 online).

To assess the prevalence of the strongest transient cell assembly activations, we computed the cumulative contribution to the epoch-wide reactivation of events with reactivation strengths up to certain values for post epoch SWS and subtracted the same measure for pre epoch SWS. This cumulative contribution (Fig. 2f) increased steadily over two orders of magnitude, and 40–50% of the net reactivation (difference between post and pre) came from events with reactivation strengths beyond the 99th percentile. Thus, rare events of elevated network synchronization or network spikes^{37,38}, although spanning only a small period of time, account for a substantial proportion of the total observed reactivation.

Preferential mPFC replay during hippocampal SPWRs

The standard systems consolidation theory holds that, following acquisition, experience-related information flows from the hippocampus toward the neocortex²⁴. Conversely, neocortical influences may contribute to selecting the pattern reactivated in the hippocampus²². We tested the relationship between mPFC cell assembly reactivation and hippocampal SPWRs, the most prominent pattern of hippocampal activation during SWS (which, similar to reactivation strength peaks, occur irregularly). Indeed, cortical assembly reactivation events occurred in concert with hippocampal SPWRs. Examining data from entire post epoch sessions (Fig. 3) revealed that virtually all reactivation peaks occurred concomitantly with a SPWR event (and also with an increase in synchronous activity of those cells with large positive weights in this signal component). We examined the ensemble spike trains corresponding to a reactivation peak for an example session (Fig. 4); at the time of the peak, virtually all cells with large positive

weights in this signal component were active (and negative weight cells reduced their activity). In this example, the two largest peaks (Fig. 4a) coincided with SPWR events (Fig. 4b) and one preceded a delta wave (Fig. 4c). The ensemble spike trains corresponding to a reactivation peak are shown (Fig. 4d,e); at the time of the peak, virtually all cells with large positive weights in this signal component were active.

The average reactivation strength in post epoch SWS was considerably greater for bins coinciding with SPWRs than for non-SPWR bins (all comparisons $P < 0.005$, $n = 8$, 37 and 225, including only sessions with reliable

discrimination of ripple signals; Fig. 5a). The effect was stronger for components with higher encoding strengths (Pearson's correlation test, $P < 10^{-20}$; Supplementary Fig. 7). The SPWR-triggered average (Fig. 5b) showed that the reactivation strength for signal components during post, but not pre, epoch SWS increased by ~70% at the time of the sharp waves with respect to baseline ($P < 10^{-10}$, $n = 270$). Reactivation in mPFC declined to baseline values within 1 s before and after the peak of the SPWR events. No such effect was found for

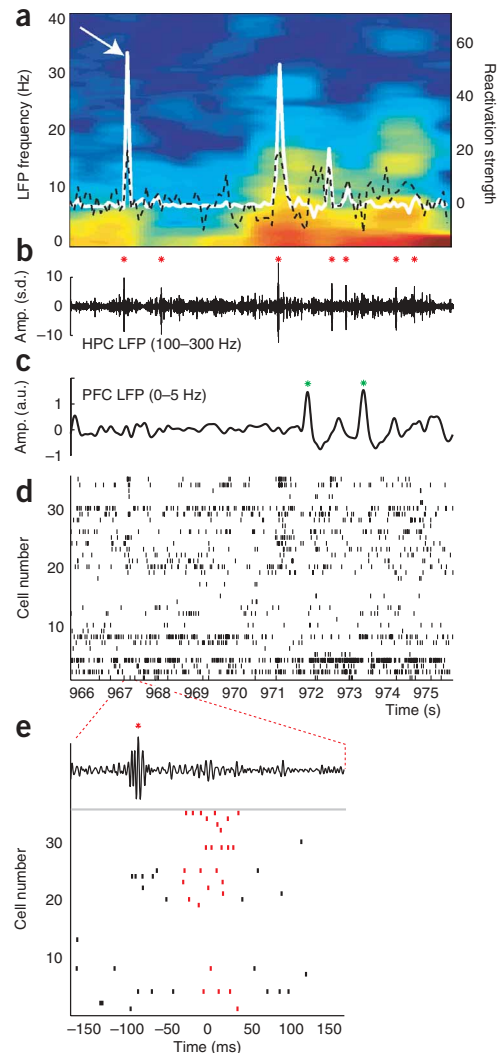


Figure 4 Example of reactivation strength peaks coinciding with hippocampal SPWR. (a) Reactivation strength (white traces, right axis) of the signal component superimposed on the mPFC LFP spectrogram (left axis). This is expanded from the zone shown in Figure 3a (post epoch), delimited by the dashed line. The black dashed line represents the normalized population firing rate. (b) The bandpass-filtered hippocampal LFP (100–300 Hz) shows ripple events (red asterisks); the signal is normalized by its s.d. (c) Bandpass-filtered (0–5 Hz) PFC LFP. Delta waves are denoted by green asterisks. (d) Raster plot of spike trains from the PFC cells sorted by principal component weight magnitude (as in Fig. 1a). (e) Expansion of the 300 ms surrounding the peak indicated by an arrow in a. Red rasters represent spikes occurring in the bin of peak reactivation strength. This example also shows two delta waves in the cortical LFP, with preceding and following increases in population activity (putative UP states⁴¹). See Supplementary Figure 7 for statistical analysis of reactivation/slow oscillation interactions.

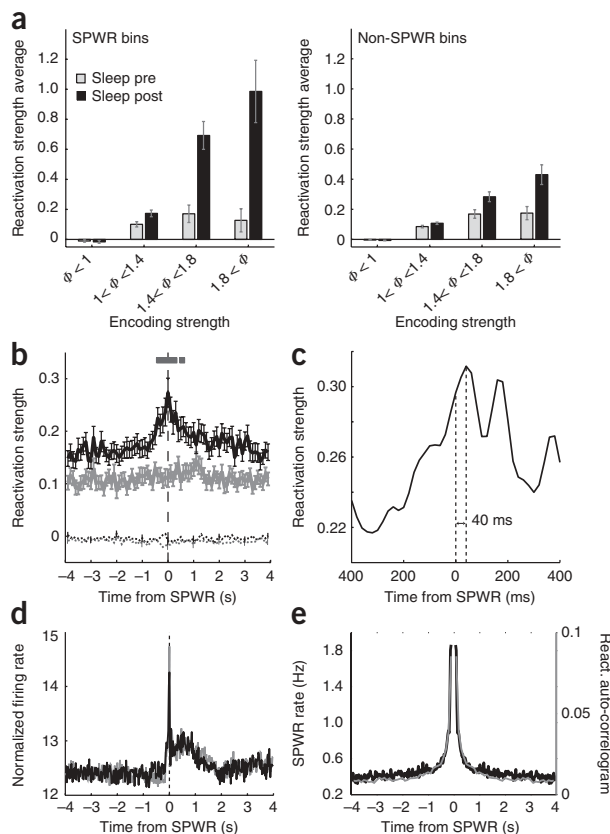


Figure 5 Prefrontal memory replay is enhanced during hippocampal SPWRs. **(a)** Average reactivation strength for nonsignal ($\phi < 1$) and signal components (grouped by encoding strength) during pre and post epoch SWS for SPWR bins (left) and non-SPWR bins (right). **(b)** Event-triggered average (ETA) of reactivation strength centered on hippocampal SPWRs for all analyzed signal components during pre SWS (gray) and post SWS (black) and for nonsignal components (black and gray, respectively). We observed an increase around SPWRs during the post epoch for signal components only. **(c)** Expanded view of the central portion of **b** for post SWS (time bins of 20 ms) with an increased reactivation peak 40 ms after SPWRs. **(d)** ETA of spiking probability density of multi-unit activity relative to SPWR occurrences, averaged over all recording sessions. Gray, pre SWS; black, post SWS. No difference was observed between the pre and post epochs. **(e)** Autocorrelogram for post SWS reactivation strengths (red) and SPWR occurrences (black) showing similar decay time constants.

oscillatory activity in the delta (2–4 Hz) and spindle (10–20 Hz) ranges ($P < 10^{-5}$ for all, t test; **Supplementary Fig. 5**). Therefore, we tested the correlation between reactivation strength and LFP markers of slow oscillations. First, we considered delta waves; that is, large positivities of the depth cortical LFP associated with states of reduced cortical activity (DOWN states) and with the K-complex phase characterized by absence of spindles⁴⁰. During post epoch SWS, the reactivation strength for signal components showed a significant ($P < 0.001$, t test) increase ~ 400 ms before the peak of the delta wave (**Fig. 6a**). This was experience dependent and possibly memory related, as the modulation was smaller for pre epoch SWS and null for nonsignal components. The timing of hippocampal SPWRs relative to delta peaks closely resembled that of mPFC reactivation (**Fig. 6a**). In contrast, mPFC ensemble activity showed a different profile (**Fig. 6a**), with a minimum immediately before the peak of the delta wave, but symmetric peaks before and after (with a return to baseline in 500–1,000 ms). The second peak was not associated with an increase in reactivation.

To further investigate this relationship, the same analysis was carried out with reference to putative DOWN to UP state transitions (putative DOWN states were defined as a decrease of neural activity in windows of at least 80 ms; **Fig. 6b**). The relationships between reactivation and SPWR occurrence with these transitions were comparable with our delta wave results. Because spindles (bouts of 10–20 Hz oscillations) appear at the onset of UP states⁴¹, we examined their correlation with reactivation strength. In general, signal reactivation tended to occur before spindle episodes. Reactivation event-triggered averages centered on spindles troughs were asymmetric, with an increase in reactivation in the ~ 1 s preceding spindles compared with the period thereafter ($P < 0.001$, t test; **Fig. 6c**). As was the case for delta waves, the increased pre-spindle reactivation over a broad time scale echoed the increased probability for hippocampal sharp waves preceding spindles (**Fig. 6c**). In contrast, the population activity modulation showed a symmetrical time course that peaked at the time of spindle events (**Fig. 6c**). The respective cross-correlograms and event-time averages of reactivation relative to these three cortical events were strongly correlated (**Fig. 6d**). Thus, coupling between reactivation and sharp waves primarily structured the relationship between the reactivation time course and cortical slow oscillations.

Replay of ‘choice point’ activity increases with learning

The principal component analysis that characterized the time course of experience-related pattern reactivation during sleep may conversely be employed to find out which aspects of the neural assembly coactivation during task performance are replayed during sleep. For this, we computed principal components from the ensemble neural activity during sleep pre and post epochs, separately for SPWR and inter-SPWR time bins. Those templates were matched to the activity during

nonsignal components. A similar analysis at higher time resolution (**Fig. 5c**) showed that reactivation peaked ~ 40 ms after SPWR occurrences, which is compatible with the transmission delay measured for prefrontal responses to hippocampal stimulation³⁹ (the second peak in the event triggered histogram is probably a result of the frequent occurrence of sharp wave ‘doublets’). On the other hand, overall ensemble mPFC activity (of all recorded neurons, including those not involved in signal components) showed a qualitatively different, sharply asymmetric profile with respect to SPWR occurrences (**Fig. 5d**); on average, prefrontal population activity transiently increased with the SPWRs and maintained sustained activity thereafter^{20,21} (with no difference between pre and post epochs). This sustained post-SPWR activity contrasts with the faster decay of signal reactivation, arguing against an explanation of the latter solely in terms of general population activity fluctuations. Furthermore, autocorrelograms of both reactivation strength and SPWR occurrences (**Fig. 5e**) decayed with very similar time constants (150 and 160 ms for exponential fits, respectively), suggesting that the clustering in time of SPWRs is reflected by a similar grouping of mPFC reactivation events.

Relation of slow oscillations with SPWRs and mPFC replay

A hallmark of cortical activity during SWS is slow oscillations²⁶, which trigger and orchestrate LFP waves in the delta (2–4 Hz) and spindle (10–20 Hz) ranges. Reactivation episodes in the hippocampus and neocortex occur during the slow oscillation phase with high neural activity¹⁶ (UP state) and are correlated with hippocampal SPWRs, but little is known about the precise temporal relation between cortical oscillatory phenomena, hippocampal activity and neocortical reactivation. When we examined the relationship between mPFC reactivation and SWS oscillations (**Fig. 2a**), we found that episodes of strong replay were significantly concentrated into periods of elevated prefrontal LFP

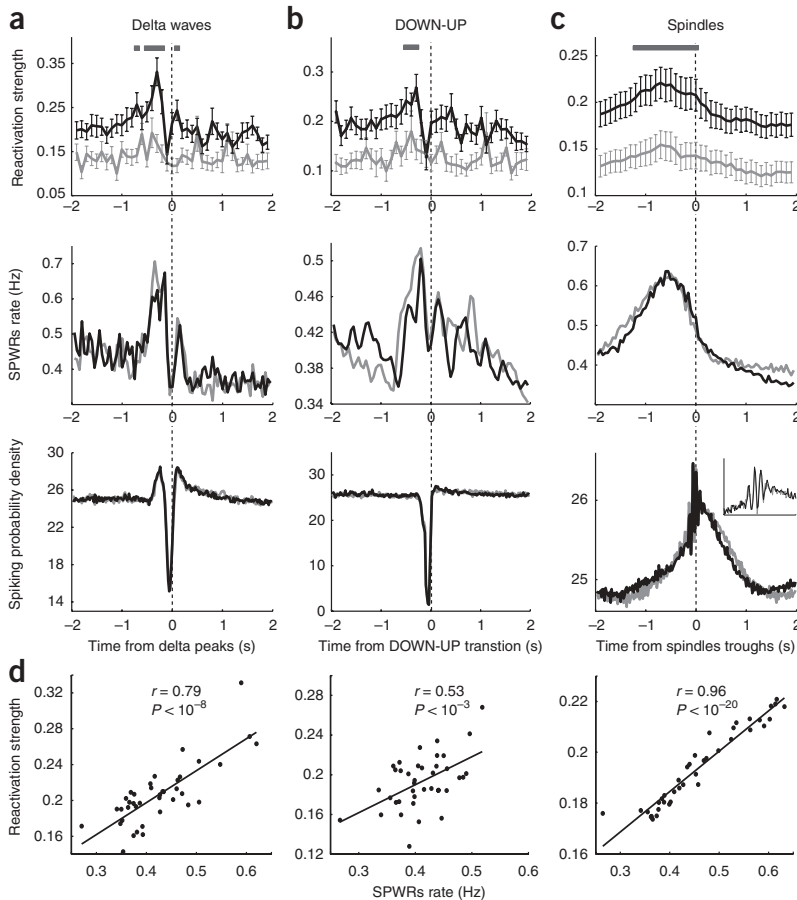


Figure 6 Reactivation strength relative to mPFC LFP events. **(a)** Cortical delta wave peaks. Top, ETA of signal components' reactivation strength centered on cortical delta wave peaks for pre (blue) and post epoch SWS (red). Gray bars indicate significantly ($P < 0.001$, t test) higher reactivation strengths for signal components during post SWS with respect to baseline (defined as the average reactivation strength from -4 to -2 s and from 2 to 4 s from the delta wave peak). Reactivation had a significant peak preceding the delta wave. Middle, cross-correlogram of SPWR occurrences relative to delta peaks. SPWRs tended to occur more frequently just before delta peaks, similar to reactivation. Bottom, spiking probability density of multi-unit activity relative to delta waves and averaged over all recording sessions. Gray, pre SWS; black, post SWS. Prefrontal cells showed a strong decrease in firing at the time of the delta peak, preceded and followed by activity increases. **(b)** Same plot as in **a**, but centered on putative DOWN to UP state transitions (as defined by population average firing rate). Results are comparable with those shown in **a** except for spiking probability, which only showed a marked deflection during the DOWN state. **(c)** Same plots as in **a**, but centered on spindle troughs. Top, reactivation strength was significantly higher ($P < 0.05$) for over 1 s before spindles. Middle, hippocampal SPWR ETA centered on spindles showed a similar profile. Bottom, spike activity increased at the time of spindles. **(d)** Comparison between time relationship of reactivation strength and SPWR occurrence relative to the three cortical events presented in **a–c** showing a high correlation in each case.

the awake period. For a significant number of components ($P < 0.05$) extracted from SPWR bins (Fig. 7a,b, Supplementary Table 1 and Supplementary Fig. 12 online), coactivations became stronger as the rat started the first run of correct trials marking rule acquisition. This difference was not significant ($P > 0.05$) for principal components computed from the inter-SPWR intervals in post or pre epochs. This was not simply a result of the elapsed time during the session, as there was no such difference between the two halves of those sessions where no rule learning occurred.

Furthermore, the principal components computed from post SPWR appeared primarily when the rat was on the central platform of the Y maze: that is, the point at which it was required to select the behavioral response (Fig. 7a,c). A significant effect ($P < 0.05$) of learning on the spatial distribution of principal components from post SPWR appeared only in the part of the maze going from the central platform toward the end of the target arm (Fig. 7c). Moreover, a factor analysis of these spatial distributions revealed that the two most important factors were concentrated on the platform and on the target arm, respectively (Fig. 7d and Online Methods).

Rule acquisition was not accompanied by a change in the principal measures of the rat behavior, including trial duration, trajectories (which followed the same stereotyped paths before and after rule acquisition) and running speed at each point of the trajectory (Supplementary Fig. 12). Moreover, the greater contribution to reactivated patterns during SPWRs from trials occurring immediately after rule acquisition is not likely to result from changes in the general sensory experience other than reward. To test this hypothesis, we compared trials occurring before and after spontaneous strategy shifts effected by the rat that did not lead to acquisition of the rewarded strategy in

sessions where no learning took place (rats made such shifts while seeking the correct rule by trial and error). In these cases, no difference was observed in contribution to reactivated patterns during SPWRs (see Online Methods, Fig. 7b and Supplementary Fig. 13 online).

DISCUSSION

We found new links between learning and the dynamics of replay in the mPFC and hippocampus and our study has three main results. First, mPFC replay occurred in transient episodes, corresponding to the activation of distinct cell assemblies. Second, prefrontal replay, although it occurred throughout the post-training SWS, mostly coincided with hippocampal sharp wave events and therefore with increased hippocampal replay and hippocampal/neocortical interactions. Third, and most importantly, in sessions with rule learning, mPFC replay during hippocampal sharp waves principally concerned neural activity patterns emerging only after the acquisition.

These results demonstrate the relationship between memory replay, the task phase where the activity patterns originated and task performance level in a dynamic setting, wherein the rats were obliged to continually adapt to new rules. Notably, when the rat adhered to the newly acquired rule, the patterns that contributed the most to memory replay in mPFC during hippocampal SPWRs were those appearing at the point at which the rat committed to choosing a goal arm. During the trial, the contribution to hippocampal-related replay in the mPFC climbed steadily on the departure arm (where activity may predict the choice⁴²), to peak at the arms intersection. Furthermore, preferentially replayed patterns arose just when rats began a series of correct trials at criterion with respect to the new rule. Prefrontal cortical activity reflecting newly learned associations has been found to emerge only

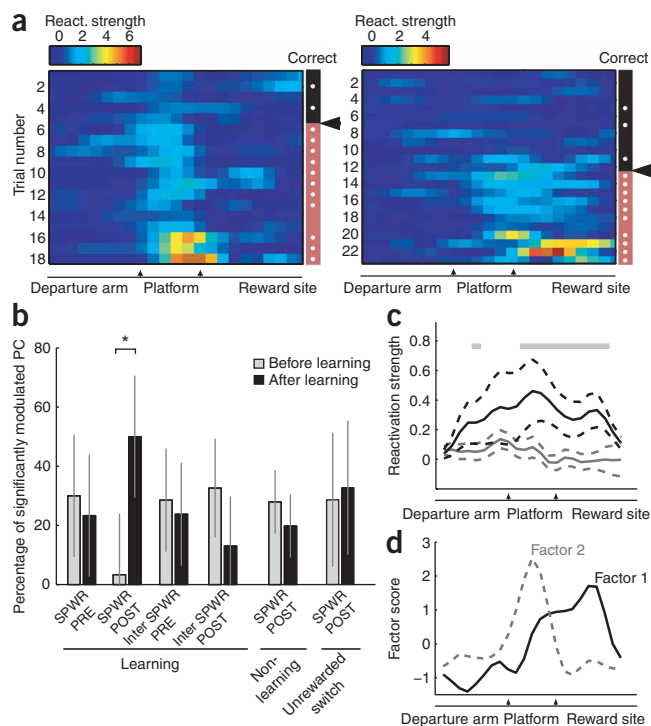


Figure 7 Replayed activity during SPWRs is correlated to rule acquisition. **(a)** The activation strength of principal components computed from post epoch SPWR activity is shown in peri-event color rasters for two examples of sessions with rule learning (in different rats). The maze was linearized and divided into 25 equal bins. Rewarded trials are marked with white dots at right. The arrows indicate the learning point (see Online Methods). The vertical bar to the right of each display shows the rat behavior: white dots correspond to the rewarded trial. The black arrow denotes the point at which a strategy shift towards the rewarded rule was detected (defined in Online Methods). **(b)** Incidences of principal components computed from pre or post epoch SPWRs or inter-SPWR epochs during sessions with learning ($n = 10$, four rats). The black bars are data recorded after learning occurred in the session (above arrows in **a**), whereas gray bars are from before learning. Significantly more post SPWR activity patterns were positively correlated with activity after learning than with activity before learning ($P < 0.05$, two-way ANOVA followed by t test). No such difference was observed for pre SPWR patterns or for inter-SPWR patterns, or when a comparison was made between the first and second half of nonlearning sessions or when the rat switched between two erroneous, unrewarded strategies. **(c)** Average spatial distribution of all fifteen principal components computed from post epoch SPWR that were significantly higher on average after learning (gray trace) than before (black trace). The gray bar (above) marks significantly different bins ($P < 0.05$, t test, $n = 10$). Dashed lines indicate s.e.m. **(d)** Factor analysis scores of data from **c** showed two profiles with peaks at the goal arms and the central choice point (see Online Methods).

after the associations have been acquired³⁴. The time course that we observed for replay is similar to that seen in the latter study³⁴: coactivations preferentially replayed in SPWRs arose after the rat started to perform at criterion. These effects were probably not a consequence of different neural activity statistics at the choice point, as they remain specific for post SPWR patterns and are not found for patterns extracted from the pre epoch or from inter-SPWR intervals.

The hippocampal involvement in the formation of these cell assemblies is likely to be critical; notably, task events linked to preferential replay were also marked by increased coherence between the hippocampus and mPFC LFP in the theta range (Benchenane *et al.*, *Soc. Neurosci. Abstr.* **690.15**, 2008). Thus, on learning, the choice point is the site of increased hippocampal-prefrontal coherence. Cell assemblies activated at those times are prominently replayed in the mPFC during SPWRs, that is, when the hippocampal-neocortical interaction is at its peak^{19–22}. Coactivations may reflect prefrontal cell assemblies representing new information. The nature of this new information could be twofold: on one hand, it could represent an emerging representation of the newly learned rule, and on the other hand, it might reflect processes that take place after learning has occurred (for example, a representation that is activated by a consistent stream of reward). The reward signal may ‘tag’ the representations, making them more likely to be replayed. In support of this hypothesis, we analyzed those periods in which rats were searching for the correct rule by trial and error. We found that when the rat makes spontaneous strategy shifts to unrewarded rules, patterns from before or after the shift were indifferently replayed in sleep (Supplementary Fig. 13). Thus, the consequent changes in the general sensory experience do not cause the resulting neural patterns to be replayed more or less strongly. Also, this seems to rule out interpretation of our results in terms of repetitive experience resulting in stronger replay¹⁸. Because preferential replay only occurs during SPWRs, we speculate that regulation of hippocampal-prefrontal interactions could result from dopaminergic, reward-related signals⁴³. The replay of reward tagged patterns, in concert

with the hippocampal replay during SPWRs^{14,44}, would mark the initial period of system-wide consolidation.

The high temporal resolution of our coactivation measures revealed another feature of prefrontal replay: its detailed time course. Replay was largely accounted for by brief events, with durations on the order of 100 ms or less (for example, see Fig. 2). In each event (similar to network spikes; see Supplementary Discussion online), a substantial number of cells were coactivated. The observed, highly irregular time course of reactivation strength suggests that replay is not simply a result of changes in the probability of co-firing for cell pairs (a straightforward consequence of a change in the efficacy of the connection between the two cells¹⁸), but is probably generated by global network effects induced, for example, by the excitatory feedback connections in the mPFC.

Transient replay is concentrated during SWS (Supplementary Fig. 5). This differs from studies of hippocampal reactivation that showed nonzero replay during restful, nonsleep periods¹⁴. This may be an intrinsic difference between the hippocampus and the mPFC or it could be a result of a different sensitivity of our analysis methods. In SWS, replay occurred principally in proximity to hippocampal SPWRs (Figs. 3–5). Thus, besides being important for the hippocampal replay of newly formed memories¹⁴, SPWRs also coincided with an increase of overall reactivation in hippocampal output structures. Moreover, SPWRs are related to increases in cortical activity and transitions to UP states^{20,21}. It is therefore possible that the correlation between cortical replay and SPWRs that we observed may be a result of a hippocampal influence on mPFC. At present, though, it is difficult to speculate whether reactivation events originate in the hippocampus or the neocortex; our data indicate that prefrontal replay peaks ~40 ms after SPWR occurrences, a latency that is compatible with observations of prefrontal responses to hippocampal stimulation³⁹ and thus with hippocampus triggering this extra-hippocampal replay. On the other hand, it is known that the neocortical slow oscillations⁴⁵ influence the membrane potential of hippocampal cells²³ and the probability of occurrence of SPWR²²; such an influence would allow the neocortex to contribute to selecting the information reactivated in the hippocampus. As a third alternative, the concurrent prefrontal cortical replay and hippocampal SPWRs may be the result of particular neocortical network states that simultaneously regulate when hippocampal reactivation takes

place and facilitate the transfer of reactivated patterns between brain regions²⁰. In a more radical version of this hypothesis, the representation of a memory item could integrally involve both hippocampus and neocortex by the time of initial encoding, and thereafter this would reactivate as a whole and manifest itself simultaneously in both structures. In fact, analyses of data from our experiments (Benchenane *et al.*, *Soc. Neurosci. Abstr.* **690.15**, 2008) show that signs of hippocampal/neocortical interplay were already present when the replayed representation was probably formed, in the form of theta-band coherence.

In conclusion, the acquisition of new rules involves the hippocampocortical network; during the ensuing sleep, the PFC activity patterns during hippocampal SPWRs reflect the neural patterns that occurred during the task phase, particularly when a rule has been learned, just at the time when hippocampocortical coherence is enhanced. Whether the predominating causal influence in this dialog is cortical or hippocampal, or rather this corresponds to an emergent system-wide representation of information, our results show a possible mechanism by which task-relevant learned information can be expressed and reactivated in the prefrontal cortex. This would be contingent to SPWRs for newly formed memories, but more uncoupled from the hippocampus for more distant memory traces.

METHODS

Methods and any associated references are available in the online version of the paper at <http://www.nature.com/natureneuroscience/>.

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We thank P. Tierney for valuable discussions and help with the surgical procedures, J.-M. Deniau, Y. Gioanni, A.-M. Thierry, M.B. Zugaro, M. Cencini and A. Aubry for interesting discussions, S. Doutremere for histology, D. Hopkins and N. Quenec'h du for the anatomical reconstructions, V. Douchamps for help with the experiments, K. Gothard, K. Hoffman and A. Treves for critical readings of an earlier version of the manuscript, and A. Berthoz for support throughout the project. This work was supported by Fondation Fyssen (F.P.B.), Fondation pour la Recherche Medicale (A.P.), and European Community contracts FP6-IST 027819 (Integrating Cognition, Emotion and Autonomy), FP6-IST-027140 (Bayesian Approach to Cognitive Systems) and FP6-IST-027017 (NeuroProbes).

AUTHOR CONTRIBUTIONS

S.I.W., F.P.B. and M.K. designed the experiment, F.P.B., M.K. and A.P. performed the experiments, A.P., F.P.B. and K.B. designed the analysis techniques, A.P. analyzed the data, and F.P.B., A.P. and S.I.W. wrote the paper.

Published online at <http://www.nature.com/natureneuroscience/>

Reprints and permissions information is available online at <http://www.nature.com/reprintsandpermissions/>

- Malenka, R.C. & Nicoll, R.A. Long-term potentiation—a decade of progress? *Science* **285**, 1870–1874 (1999).
- Marr, D. Simple memory: a theory for archicortex. *Phil. Trans. R. Soc. Lond. B* **262**, 23–81 (1971).
- McClelland, J.L., McNaughton, B.L. & O'Reilly, R.C. 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 (1995).
- Dudai, Y. The neurobiology of consolidations or how stable is the engram? *Annu. Rev. Psychol.* **55**, 51–86 (2004).
- Squire, L.R. Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* **99**, 195–231 (1992).
- Moscovitch, M., Nadel, L., Winocur, G., Gilboa, A. & Rosenbaum, R.S. The cognitive neuroscience of remote episodic, semantic and spatial memory. *Curr. Opin. Neurobiol.* **16**, 179–190 (2006).
- Winocur, G., Moscovitch, M. & Sekeres, M. Memory consolidation or transformation: context manipulation and hippocampal representations of memory. *Nat. Neurosci.* **10**, 555–557 (2007).
- Tse, D. *et al.* Schemas and memory consolidation. *Science* **316**, 76–82 (2007).
- Marshall, L., Helgadottir, H., Mölle, M. & Born, J. Boosting slow oscillations during sleep potentiates memory. *Nature* **444**, 610–613 (2006).
- Walker, M.P. & Stickgold, R. Sleep, memory, and plasticity. *Annu. Rev. Psychol.* **57**, 139–166 (2006).
- Wagner, U., Gais, S., Haider, H., Verleger, R. & Born, J. Sleep inspires insight. *Nature* **427**, 352–355 (2004).
- Wilson, M.A. & McNaughton, B.L. Reactivation of hippocampal ensemble memories during sleep. *Science* **265**, 676–679 (1994).
- Nádasdy, Z., Hirase, H., Czurko, A., Csicsvari, J. & Buzsáki, G. Replay and time compression of recurring spike sequences in the hippocampus. *J. Neurosci.* **19**, 9497–9507 (1999).
- Kudrimoti, H.S., Barnes, C.A. & McNaughton, B.L. Reactivation of hippocampal cell assemblies: effects of behavioral state, experience and EEG dynamics. *J. Neurosci.* **19**, 4090–4101 (1999).
- Hoffman, K.L. & McNaughton, B.L. Coordinated reactivation of distributed memory traces in primate neocortex. *Science* **297**, 2070–2073 (2002).
- Ji, D. & Wilson, M.A. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat. Neurosci.* **10**, 100–107 (2007).
- Euston, D.R., Tatsuno, M. & McNaughton, B.L. Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science* **318**, 1147–1150 (2007).
- O'Neill, J., Senior, T.J., Allen, K., Huxter, J.R. & Csicsvari, J. Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nat. Neurosci.* **11**, 209–215 (2008).
- Siapas, A.G. & Wilson, M.A. Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron* **21**, 1123–1128 (1998).
- Sirota, A., Csicsvari, J., Buhl, D. & Buzsáki, G. Communication between neocortex and hippocampus during sleep in rodents. *Proc. Natl. Acad. Sci. USA* **100**, 2065–2069 (2003).
- Battaglia, F.P., Sutherland, G.R. & McNaughton, B.L. Hippocampal sharp wave bursts coincide with neocortical 'up-state' transitions. *Learn. Mem.* **11**, 697–704 (2004).
- Isomura, Y. *et al.* Integration and segregation of activity in entorhinal-hippocampal subregions by neocortical slow oscillations. *Neuron* **52**, 871–882 (2006).
- Hahn, T.T., Sakmann, B. & Mehta, M.R. Phase-locking of hippocampal interneurons' membrane potential to neocortical up-down states. *Nat. Neurosci.* **9**, 1359–1361 (2006).
- Buzsáki, G. Two-stage model of memory trace formation: a role for 'noisy' brain states. *Neuroscience* **31**, 551–570 (1989).
- Buzsáki, G. Hippocampal sharp waves: their origin and significance. *Brain Res.* **398**, 242–252 (1986).
- Steriade, M., Nuñez, A. & Amzica, F. A novel slow (<1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. *J. Neurosci.* **13**, 3252–3265 (1993).
- Frankland, P.W. & Bontempi, B. The organization of recent and remote memories. *Nat. Rev. Neurosci.* **6**, 119–130 (2005).
- Gais, S. *et al.* Sleep transforms the cerebral trace of declarative memories. *Proc. Natl. Acad. Sci. USA* **104**, 18778–18783 (2007).
- Jay, T.M. & Witter, M.P. Distribution of hippocampal CA1 and subicular efferents in the prefrontal cortex of the rat studied by means of anterograde transport of *Phaseolus vulgaris* leucoagglutinin. *J. Comp. Neurol.* **313**, 574–586 (1991).
- Jay, T.M., Burette, F. & Laroche, S. NMDA receptor-dependent long-term potentiation in the hippocampal afferent fibre system to the prefrontal cortex in the rat. *Eur. J. Neurosci.* **7**, 247–250 (1995).
- Takashima, A. *et al.* Declarative memory consolidation in humans: a prospective functional magnetic resonance imaging study. *Proc. Natl. Acad. Sci. USA* **103**, 756–761 (2006).
- Floresco, S.B., Braakman, D.N. & Phillips, A.G. Thalamic-cortical-striatal circuitry subserves working memory during delayed responding on a radial arm maze. *J. Neurosci.* **19**, 11061–11071 (1999).
- Jones, M.W. & Wilson, M.A. Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLoS Biol.* **3**, e402 (2005).
- Pasupathy, A. & Miller, E.K. Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature* **433**, 873–876 (2005).
- Baeg, E.H. *et al.* Learning-induced enduring changes in functional connectivity among prefrontal cortical neurons. *J. Neurosci.* **27**, 909–918 (2007).
- Birrell, J.M. & Brown, V.J. Medial frontal cortex mediates perceptual attentional set shifting in the rat. *J. Neurosci.* **20**, 4320–4324 (2000).
- Beggs, J.M. & Plenz, D. Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures. *J. Neurosci.* **24**, 5216–5229 (2004).
- Mongillo, G., Barak, O. & Tsodyks, M. Synaptic theory of working memory. *Science* **319**, 1543–1546 (2008).
- Dégenétais, E., Thierry, A.M., Glowinski, J. & Gioanni, Y. Synaptic influence of hippocampus on pyramidal cells of the rat prefrontal cortex: an *in vivo* intracellular recording study. *Cereb. Cortex* **13**, 782–792 (2003).
- Amzica, F. & Steriade, M. Electrophysiological correlates of sleep delta waves. *Electroencephalogr. Clin. Neurophysiol.* **107**, 69–83 (1998).
- Steriade, M. Neuronal substrates of spike-wave seizures and hypsarrhythmia in corticothalamic systems. *Adv. Neurol.* **97**, 149–154 (2006).
- Baeg, E.H. *et al.* Dynamics of population code for working memory in the prefrontal cortex. *Neuron* **40**, 177–188 (2003).
- Seamans, J.K., Floresco, S.B. & Phillips, A.G. D1 receptor modulation of hippocampal-prefrontal cortical circuits integrating spatial memory with executive functions in the rat. *J. Neurosci.* **18**, 1613–1621 (1998).
- Lee, I. & Kesner, R.P. Differential contribution of NMDA receptors in hippocampal subregions to spatial working memory. *Nat. Neurosci.* **5**, 162–168 (2002).
- Steriade, M. Grouping of brain rhythms in corticothalamic systems. *Neuroscience* **137**, 1087–1106 (2006).

ONLINE METHODS

Animals. Four Long-Evans (pigmented) male rats (R. Janvier, Le Genest-St-Isle, France) weighing 250–300 g at arrival were handled daily. All experiments were carried out in accordance with institutional (CNRS Comité Opérationnel pour l'Éthique dans les Sciences de la Vie) and international (US National Institute of Health guidelines) standards and legal regulations (Certificate no. 7186, French Ministère de l'Agriculture et de la Pêche) regarding the use and care of animals.

After habituation to the experimental environment, rats were anesthetized with intramuscular xylazine (Rompun, 0.1 ml) and intraperitoneal pentobarbital (35 mg per kg of body weight). A drive containing seven tetrodes (six recording, plus one reference) was implanted through the skull above the right medial PFC (anterior-posterior, 3.5–5 mm; medial-lateral, 0.5–1.5 mm). Each tetrode was contained in a 30 gauge hypodermic tube, with the tubes arranged in two parallel, adjacent rows. Tetrodes were twisted bundles of polyimide-coated nichrome wire (13 μ m in diameter, Kanthal); the drive allowed independent adjustment of tetrode depth. After dura retraction, the rows of cannulae were implanted parallel to the sagittal sinus so that they targeted the superficial and deep layers of the medial bank of the cortex. A separate micro-drive containing three tetrodes was targeted to the ventral hippocampus (anterior-posterior, –5.0 mm; medial-lateral, 5.0 mm). Each tetrode was electrically connected in a single-electrode configuration (all channels shorted together) and used for LFP recordings. For these, a screw implanted on the occipital bone above the cerebellum served as the reference. The hippocampal tetrodes were lowered to the CA1 pyramidal layer; the depth was adjusted with the help of LFP signs (flat sharp waves, strong ripple oscillations). After surgery, rats recovered for at least 2 weeks while the tetrodes were lowered to reach the prelimbic area (main drive) and the CA1 pyramidal layer (hippocampal micro-drive). Between sessions, tetrodes were gradually lowered to probe different dorso-ventral levels in the prelimbic area.

Apparatus and pre-training. An elevated, wooden Y maze was coated with waterproof black paint. The three arms were 85 cm long, 8 cm wide and separated by 120 degrees. The arms connected on a circular central platform, which could be elevated to block the access to all arms. A wooden box at the end of each arm contained a reward well and a small light bulb that illuminated the box and the end portion of the arm. Solenoids controlled the delivery of drops of a liquid reward (chocolate- or strawberry-flavored milk). Infrared photo-detectors, at the entrance of each arm and at the reward wells, detected rat access to the arms and arrival at the reward site. The solenoids, light bulbs and photo-detectors were connected to a micro-controller built into the data acquisition system (Power1401, CED). The micro-controller managed reward delivery and arm illumination on the basis of the behavioral task logic. A monochrome camera was placed above the maze and tracked rat positions. The video signal was acquired and synchronized with behavioral and electrophysiological data by the data acquisition software (Spike2, CED). The position of LEDs placed on the recording headstage was tracked offline by dedicated software (MaxTRAQ, Innovision Systems).

Behavioral task. Rats performed an attentional set shift task on a Y maze. Such extradimensional set shift tasks have been shown to require the function of the mPFC in rats³⁶. This parallels the involvement of the human PFCs in the Wisconsin card sorting task, which inspired our experimental design. Each recording session consisted of a 20–30-min sleep or rest epoch (pre epoch), in which the rat remained undisturbed in a padded flowerpot placed on the central platform of the maze, followed by an awake epoch, in which the rat performed the behavioral task described below for 20–40 min, and then by a second 20–30-min sleep or rest epoch (post epoch; same situation as in pre epoch).

The first recording sessions were the first time that the rats encountered the behavioral task. Rats started each trial in the same arm (the departure arm). One of the two other (choice) arms was illuminated at random (pseudo-random schedule; runs of more than four consecutive trials with the same illuminated arm were avoided, as were repeated bouts of imposed alternation between the two arms). After that, the central platform was lowered, allowing the rat to access the choice arms.

Only one of the choice arms was rewarded, according to one of four contingency rules. Two contingency rules were spatially guided (always go to the right arm or to the left arm), the other two were cue guided (go to the illuminated arm or to the dark arm). The animal had to learn each rule by trial and error, as it was not signaled by any cue. Once the rat reached a criterion of ten consecutive correct trials or one error out of 12 trials, the rule was changed with no further cue to the rat. Rule changes were extra-dimensional, from a spatially guided rule to a cue-guided rule, and vice versa. The learning point was defined as the first trial of a block of at least three consecutive rewarded trials after which the performance until the end of the session was above 80%. Because of an operator mistake, an intradimensional shift was applied once. Typically, the trials necessary for acquisition of a rule spanned more than one session; hence rule shifts did not occur in all sessions.

All four rats learned the right and light arm rules (according to above criteria), whereas only two learned the left arm rule and one learned the dark arm rule.

Electrophysiology. Tetrode recordings were obtained from up to six tetrodes implanted in prelimbic area and hippocampal LFPs were obtained from tetrodes implanted in the CA1 pyramidal layer in the mid-ventral part of hippocampus that projects to mPFC. Because tetrodes were not moved after each recording session, it is possible that the same cells were recorded more than once, usually in different behavioral situations (current rule, performance level, etc.).

For electrophysiological recordings, signals were fed into a unit-gain headstage pre-amplifier (HS-54, Neuralynx) and then, through a tether cable, to amplifiers (Lynx-8, Neuralynx), where all signals were amplified 2,000-fold. Single-unit recording signals were bandpass filtered between 600 and 6,000 Hz, whereas LFP signals were filtered between 0.1 and 475 Hz (cortex) and 1 and 475 Hz (hippocampus). Data were digitized and stored by a Power1401 (CED) acquisition system, which was controlled by Spike2 software. Single-unit data were sampled at 25 kHz and a 1.3-ms sample was time-stamped and stored for all the channels in a tetrode whenever any of the four exceeded a pre-set threshold. LFPs were sampled and stored at 2 kHz.

Histology. At the end of the experiments, a small electrolytic lesion was made by passing a small cathodal direct current (20 μ A, 10 s) through each recording tetrode to mark their tip locations. The rats were then deeply anesthetized with pentobarbital. Intracardial perfusion with saline was followed by 10% formalin saline (vol/vol). Histological sections were stained with cresyl violet. The electrode tracks were reconstructed, verifying that the recording sites were in the prelimbic area or, in exceptional cases, in the dorsal-most infralimbic cortex.

SWS detection. SWS was automatically detected. Power spectrograms of cortical and hippocampal LFPs were computed with bins of 1 s. Power in the cortical delta band (1–4 Hz), hippocampal theta (5–10 Hz), cortical spindles (10–20 Hz) and speed of head motion were clustered with a K-means algorithm. Clusters corresponding to high values of delta and spindle powers, and to a low degree of head movement, were considered as corresponding to SWS. Successive SWS clusters occurring in intervals of less than 1 min were merged and resulting time intervals of SWS smaller than 10 s were dropped.

Data preprocessing. For single-unit activity discrimination, the first three principal components of the energy-normalized waveforms were computed from spike waveforms for the four tetrodes, generating a 12-dimensional vector characterizing each spike. Those vectors were entered in the KlustaKwik⁴⁶ clustering program. The resulting classification was manually refined using the Klusters⁴⁷ interface.

For ripple detection, the hippocampal LFP from the CA1 pyramidal layer was bandpass filtered in the ripple frequency range (100–300 Hz) to detect SPWRs. The r.m.s. signal was then smoothed (20-ms Gaussian window). Only intervals for which the resulting envelope was 2 s.d. above the raw filtered signal were retained. Ripple events were time-stamped with the times of the deepest LFP troughs in these intervals if the latter were at least 5 s.d. below the r.m.s. signal baseline. Finally, SPWR discrimination criteria were considered reliable when at least 40 ripples were detected. In most sessions, at least 200 SPWRs were detected.

For spindle detection, cortical LFP signal were bandpass filtered (10–20 Hz). The r.m.s. signal was smoothed with a 100-ms Gaussian window and those time intervals 1 s.d. above the filtered signal and at intervals less than 100 ms were merged. Of the resulting intervals, only those at least 500 ms long were retained. Spindle troughs are the minima of the filtered signal during those intervals. For delta wave detection, cortical LFP signal was filtered between 0.1 and 4 Hz. Delta wave peaks were the minima at least 2 s.d. below the filtered signal.

Computational methods. Reactivation strengths describe the instantaneous replay at each time, during rest sessions, of neural coactivation patterns that were characteristic of the awake epoch. Each recording session was subdivided into three epochs (pre, awake and post). The goal of the analysis was to find the characteristic patterns of neural activity from one of these epochs (referred to as ‘template’) and to determine the extent and at which times such patterns were active during a second ‘match’ epoch. To determine when patterns from the behavioral episode were reactivated during the following sleep (the analyses of Figs. 1–6), we set the awake epoch as the template and the post epoch as the match epoch (with pre serving as control). For the analyses of Figure 7 (where we studied which phases of the behavioral task were most represented in the sleep activity), the template epoch was the post epoch (with the pre epoch as a control) and the match epoch was the awake epoch.

For all epochs, spike trains from prefrontal cells were binned in intervals of duration Δt_{bin} (100 ms unless otherwise specified). The obtained spike counts for each cell $s^i(t)$, $i = 1 \dots n$, $t = 1 \dots B$, where n is the number of cells and B is the number of time bins in the epoch, were z transformed, obtaining the Q matrix

$$Q_{it} = \frac{s^i(t) - \langle s^i \rangle}{\sigma_{s^i}}$$

where $\langle s^i \rangle = \frac{1}{B} \sum_{t=1}^B s^i(t)$ and $\sigma_{s^i} = \sqrt{\frac{1}{B-1} \left(\sum_{t=1}^B (s^i(t))^2 - \langle s^i \rangle^2 \right)}$.

The pairwise cell activity correlation matrix is then

$$C = \frac{1}{B} QQ^T \tag{1}$$

For the three behavioral epochs, we obtained three population spike count matrices, Q^{pre} , Q^{awake} and Q^{post} , and the three correlation matrices, C^{pre} , C^{awake} and C^{post} .

Previously¹², the reactivation of neural activity associated with active experience has been studied by looking at the relationship between statistical features of the activity as measured during behavior and during resting periods. To determine whether this replay possibly corresponds to a specific modification of the underlying neural circuits, we need statistical measures that are informative about the neural patterns that occurred at different moments during the experience, and not just global, epoch-wide modulation of neural activities. The simplest such measure is the reactivation of zero-lag, pairwise cell activity coactivation. To explain this at an intuitive level, suppose that cells A and B coactivate during experience more than expected by chance given their average activity. If, during subsequent sleep, the same cell pair shows increased coactivation (compared with sleep before experience), this can be taken as a sign of memory reactivation.

Previously proposed measures of reactivation of coactivation patterns¹⁴ are based on the comparison of the C matrices for the awake and post epochs, by taking a Pearson correlation coefficient of the elements of the two matrices or a partial correlation coefficient, and then examining the correlation between post and awake matrix elements once the effects of the pre-existing correlations (in the pre epoch matrix) are taken into account. This measure computes the general similarity level between the activity patterns in the awake and post epochs, but it does not provide information about the detailed time course of the activity during sleep and its instant-by-instant resemblance to experience-associated patterns. For this, we take a measure equivalent to the Pearson correlation coefficient of the elements of $C^{template}$ and C^{match} (assigning the role of template and match to two of the three behavioral epochs), except for constant terms (a graphical schematic of the method is provided in Supplementary Fig. 3):

$$M^{template-match} = \sum_{i,j,i < j} C_{ij}^{match} C_{ij}^{template} \\ = \frac{1}{2} Tr((C^{match} - I)^T (C^{template} - I))$$

$M^{template-match}$ can be decomposed into a sum over time bins during the post (pre) epoch (by using equation (1)).

$$M^{template-match} = \frac{1}{2} \sum_{i,j,i \neq j} C_{ij}^{match} C_{ij}^{template} \\ = \frac{1}{2B^{match}} \sum_{t=1}^{B^{match}} \sum_{i,j,i \neq j} Q_{it}^{match} C_{ij}^{template} Q_{jt}^{match} \\ = \frac{1}{2B^{match}} \sum_{t=1}^{B^{match}} R_0^{match}(t)$$

Thus, $C_{ij}^{template}$ can be seen as a template (technically a quadratic form) applied to the vector of multi-cell spike counts Q_{ij}^{match} at each time t during the rest epochs to produce the time series $R_0^{match}(t)$. $R_0^{match}(t)$ represents a decomposition of the epoch-wide correlation similarity into its instantaneous contributions; that is, the similarity between the current population vector at time t and the general pattern of coactivation during the template epoch. Therefore, it contains information concerning exactly when during the match epochs the patterns of coactivation are similar to those in the template. However, such a measure still combines together factors from several different cell groups, which may coactivate independently. The obtained time course may therefore be the result of averaging over these distinct patterns, which may behave quite differently from one another.

To distinguish among these cell groups (and to find which of these patterns contributed the most to the reactivation effect), we found the eigenvectors of the $C^{template}$ matrix, $p^{(l)}$ ($p^{(1)}$ is associated with the highest eigenvalue λ_1). The eigenvectors of the correlation matrix $C^{template}$ are the so-called principal components of the data $Q^{template}$ and the ones with the lowest rankings explain most of the variance of $Q^{template}$. We can then decompose the $C^{template}$ matrix into projectors, $P^{(l)}$, defined by the outer products of the eigenvectors with themselves.

$$C^{template} = \sum_l \lambda_l P^{(l)} = \sum_l \lambda_l P^{(l)} (P^{(l)})^T$$

The time series $R_0^{match}(t)$ can then be further decomposed into terms corresponding to the different eigenvectors.

$$M^{template-match} = \frac{1}{2B^{match}} \sum_l \lambda_l \sum_{t=1}^{B^{match}} \sum_{i,j,i \neq j} Q_{it}^{match} P_{ij}^{(l)} Q_{jt}^{match} \\ = \frac{1}{2B^{match}} \sum_l \lambda_l \sum_{t=1}^{B^{match}} R_l^{match}(t) \tag{2}$$

with

$$R_l^{match}(t) = \sum_{i,j,i \neq j} Q_{it}^{match} P_{ij}^{(l)} Q_{jt}^{match}$$

The time series $R_l^{match}(t)$ measures the instantaneous match of the l -th coactivation template on the ongoing activity.

By definition, the reactivation strength measure is not very sensitive to instantaneous firing rates; overall firing-rate fluctuations are canceled out because the principal components’ coefficients have a zero mean. General

firing-rate fluctuation effects on the reactivation strengths can be controlled for with a shuffling procedure (Supplementary Fig. 8).

Signal and nonsignal components. The exact distribution of singular values (root square of the eigenvalues of the correlation matrix C^{template}) of random, uncorrelated n -dimensional datasets follow the so-called Marcenko-Pastur distribution^{48,49}. In the limit $B \rightarrow \infty$ and $n \rightarrow \infty$, with $q = \frac{B}{n} \geq 1$ fixed,

$$\rho(\lambda) = \frac{q}{2\pi\sigma^2} \frac{\sqrt{(\lambda_{\max} - \lambda)(\lambda - \lambda_{\min})}}{\lambda} \quad (3)$$

where $\lambda_{\min}^{\max} = \sigma^2(1 \pm \sqrt{1/q})^2$. σ^2 is the variance of the elements of the random matrix, which is 1 here, because the Q^{template} matrix is z transformed. Equation (3) shows that the distribution vanishes for λ greater than an upper limit λ_{\max} .

Under the null hypothesis of an uncorrelated Q^{template} , the correlations between spike trains are determined only by random fluctuations and the eigenvalues of C^{template} must lie between λ_{\min} and λ_{\max} . Eigenvalues greater than λ_{\max} are therefore a sign of nonrandom correlations in the matrix, and for this reason we refer to these principal components as signal. Components associated with eigenvalues between λ_{\min} and λ_{\max} are termed nonsignal components. For the awake epoch, a total of 323 components were classified as signal and 811 were classified as nonsignal. To combine data from different sessions, we used the normalized eigenvalues (or encoding strengths) $\Phi = \lambda / \lambda_{\max}$.

Heavy-tail characterization. The epoch-wide reactivation strength (equation (2)) is given by the sum of instantaneous contributions. To analyze the contribution to this sum by the peaks making up the tail of the reactivation strength distribution, we computed the cumulative contribution

$$\langle R \rangle_{-\infty}^r = \int_{-\infty}^r uP(u)du$$

whose difference between post and pre is shown in Figure 2e. For distributions $P(u)$ with an exponential tail, this function will reach an asymptotic value, indicating that large values contribute little to the sum. Diverging values of $\langle R \rangle_{-\infty}^r$ (for example, $\propto \log(r)$) are indicative of a $P(u)$ with a tail decaying with a power law.

Detection of sleep-related patterns in activity during the task. For the analysis of Figure 7, the rest epochs (post, with pre as a control) were the templates in the analysis described above, whereas awake was the match epoch. The templates were computed separately from the bins (of constant duration 100 ms) coinciding with SPWRs and inter-SPWRs intervals. For each SPWR, we considered the bin coincident with the ripple peak and the following one, as SPWR activity. All of the remaining bins during SWS were considered inter-SPWR activity. Multinomial statistics assessed significant changes in the number of cell assemblies whose activation strength was higher after learning than before (Fig. 7b). For each epoch, a limited number of principal components were selected: for SWS, only signal components, and for SPWRs, the first three (as a result of the low dimension of the dataset available from SPWR events, significance cannot be assessed from eigenvalues as described above; three principal components appeared to be adequate to describe reactivated patterns determined from SPWRs). Comparisons of activation strengths before and after learning (with an ANOVA test) were classified in three categories: higher, lower or not significantly different. As these groups are not independent, the use of multinomial statistics was necessary to compute the 95% confidence interval of the probabilities mean of each category⁵⁰. Thus,

significantly different probabilities ($P < 0.05$) were considered as a single group using non-overlapping confidence intervals (the group data for days when a rule was acquired are summarized in Supplementary Table 1). For nonlearning days, a total of 111 principal components were considered; 31 were lower (28%) and 22 were higher (20%) when the halves of the session were compared. For nonrewarded strategy switches, 24 principal components were considered; seven were lower (29%) and eight were higher (33%).

The spatial raster plots in Figure 7 and Supplementary Figures 11 and 12 were created from a spatially linearized trajectory on the Y maze (that is, in which the right and left arms of the Y maze were combined), obtaining a path going from the departure arm, through the central platform, to the reward site at the end of the target arm (the trial-by-trial time course of the activation of post SPWR-related activity patterns is shown in Fig. 7a). The activation of post SPWR patterns increased after the strategy shift in both examples. This phenomenon was observed in all our data, as significantly ($P < 0.05$) principal components derived from post SPWR ensemble activity from all sessions increased their activation than decreased. This difference was not present for patterns extracted from pre epoch SPWRs or from pre and post epoch inter-SPWR periods, or when we examined the difference between the first and second half in sessions in which the rewarded rule was not acquired (Fig. 7b). There was a significant difference ($P < 0.05$) between the two in the portion of trajectory starting at the platform (that is, the maze choice point; Fig. 7c).

For factor analysis of when post patterns appeared in the awake epoch (Fig. 7d), we carried out a principal component analysis in which vectors were rotated to maximize the variance (the 'varimax' rotation). Two factors seemed to adequately describe the data.

Furthermore, we analyzed sessions in which the animal spontaneously switched between nonrewarded strategies. If the rat switched between more than two strategies in a session, only the last switch was considered. Blocks of trials in which the animal consistently followed a given strategy were automatically detected. Those strategies considered were right, left, light, dark or alternation (between the left and right arms, an often-observed spontaneous, but unrewarded, strategy). Because the probability that an animal performs the same behavioral strategy for six consecutive trials by chance is less than 5% (3.12%), we considered this to be the criterion to consider a block of trials as adhering to a given strategy. The criterion for a consistent strategy also accepted blocks of at least nine trials including one noncompliant trial. For example, if the animal performed a right strategy for six trials, then went to the left arm for one trial, and then went to the right arm for two trials, the whole block of nine trials was considered to be a right strategy block. Following the same principle, we allowed blocks of nine trials constituted of five compliant trials, one noncompliant trial and three compliant trials, or four compliant trials, one noncompliant trial and four compliant trials. A total of eight sessions (at least one for each rat) with such shifts were analyzed (Supplementary Fig. 13). The same analysis as in Figure 7 showed that post SPWR patterns have the same incidence for learning the first or the second rules.

46. Harris, K.D., Henze, D.A., Csicsvari, J., Hirase, H. & Buzsáki, G. Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J. Neurophysiol.* **84**, 401–414 (2000).
47. Hazan, L., Zugaro, M. & Buzsáki, G. Klusters, NeuroScope, NDManager: a free software suite for neurophysiological data processing and visualization. *J. Neurosci. Methods* **155**, 207–216 (2006).
48. Marčenko, V.A. & Pastur, L.A. Distribution of eigenvalues for some sets of random matrices. *Math USSR SB* **1**, 457–483 (1967).
49. Sengupta, A.M. & Mitra, P.P. Distributions of singular values for some random matrices. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics* **60**, 3389–3392 (1999).
50. Fitzpatrick, S. & Scott, A. Quick simultaneous confidence intervals for multinomial proportions. *J. Am. Stat. Assoc.* **82**, 875–878 (1987).