

Temporally Structured Replay of Awake Hippocampal Ensemble Activity during Rapid Eye Movement Sleep

Kenway Louie and Matthew A. Wilson*

Department of Biology
Department of Brain and Cognitive Sciences
Center for Learning and Memory
RIKEN-MIT Neuroscience Research Center
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Summary

Human dreaming occurs during rapid eye movement (REM) sleep. To investigate the structure of neural activity during REM sleep, we simultaneously recorded the activity of multiple neurons in the rat hippocampus during both sleep and awake behavior. We show that temporally sequenced ensemble firing rate patterns reflecting tens of seconds to minutes of behavioral experience are reproduced during REM episodes at an equivalent timescale. Furthermore, within such REM episodes behavior-dependent modulation of the subcortically driven theta rhythm is also reproduced. These results demonstrate that long temporal sequences of patterned multineuronal activity suggestive of episodic memory traces are reactivated during REM sleep. Such reactivation may be important for memory processing and provides a basis for the electrophysiological examination of the content of dream states.

Introduction

The hippocampus is a region of high-level sensory convergence that is crucial to the formation and encoding of memories (Zola-Morgan and Squire, 1993). Extensive work in rodents has demonstrated direct behavioral correlates of hippocampal neuronal activity, the most robust of which is the selective activation of CA1 pyramidal cells at particular locations in space (place fields) (O'Keefe and Dostrovsky, 1971). Consistent with a hippocampal role in memory encoding, these cells exhibit experience-dependent reactivation during sleep that is representative of previous behavior (Pavlides and Winson, 1989; Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996). Specifically, neurons with overlapping place fields during spatial exploration show increased coactivity during subsequent sleep. Such short-timescale mnemonic changes are associated with slow wave sleep (SWS), particularly the high-frequency ripple oscillations during which many hippocampal neurons fire in close temporal synchrony. These oscillations provide ideal physiological conditions for the Hebbian modification of synapses (Bliss and Collingridge, 1993), suggesting that SWS reactivation can drive downstream synaptic changes to encode memory representations (Buzsaki, 1989).

In contrast, the role of rapid eye movement (REM) sleep during memory consolidation is unclear. The strong association between human dreaming and REM sleep raises many questions about the information content of dream states as well as the physiological function of REM sleep. Deprivation studies demonstrate the necessity of REM sleep for the acquisition of certain types of learning (Smith, 1995), but it has been argued that REM sleep may serve a general homeostatic role rather than a specific memory-processing function (Crick and Mitchison, 1983). Although general experience-dependent changes in neural activity occur during REM sleep (Pavlides and Winson, 1989; Poe et al., 2000), efforts to detect short-timescale mnemonic activity like that observed during SWS have failed to detect such replay (Kudrimoti et al., 1999). However, unlike SWS, REM sleep is dominated by the robust theta oscillations (6–10 Hz) and EEG desynchrony that characterize the awake exploratory state, raising the possibility that reactivation during REM sleep may be temporally structured like awake neural activity.

To investigate this, we employed a behavioral task that produces distinct hippocampal firing patterns over extended durations and examined subsequent REM episodes for similar patterns of activity. Four male Long-Evans rats were chronically implanted with microelectrode arrays to record multiple single-cell activity from the CA1 region of the hippocampus (Wilson and McNaughton, 1993). Animals were trained to run along a circular track for food reinforcement, traversing three quarters of the track circumference in each trial of a four-trial sequence that was continuously repeated for the duration of the task (Figure 1A). Following acquisition of the task, electrophysiological activity was monitored during task performance (RUN) and during periods of sleep immediately before and after behavior.

Results

The ability to simultaneously record the activity of multiple neurons enables the examination of complex patterns of firing structure beyond pairwise firing biases. CA1 pyramidal cells recorded during the behavioral task displayed spiking activity that was strongly dependent upon the animal's position in space (Figure 1B). To examine the influence of mnemonic coding on hippocampal activity, analysis was restricted to pyramidal cells that were active and unambiguously isolated throughout all sleep and behavioral epochs. Consistent with previous observations of place cell activity within the hippocampus, in which ~30% of cells are typically active in any given spatial environment (Wilson and McNaughton, 1993), cells with mean RUN firing rates exceeding 0.2 Hz were identified as active, yielding ensembles consisting of between 8 and 13 simultaneously recorded neurons per session (see Experimental Procedures). While some cells were strongly modulated by location alone, other cells fired in a conjunctive manner combining both location specificity and behavioral specificity

* To whom correspondence should be addressed (e-mail: wilson@ai.mit.edu).

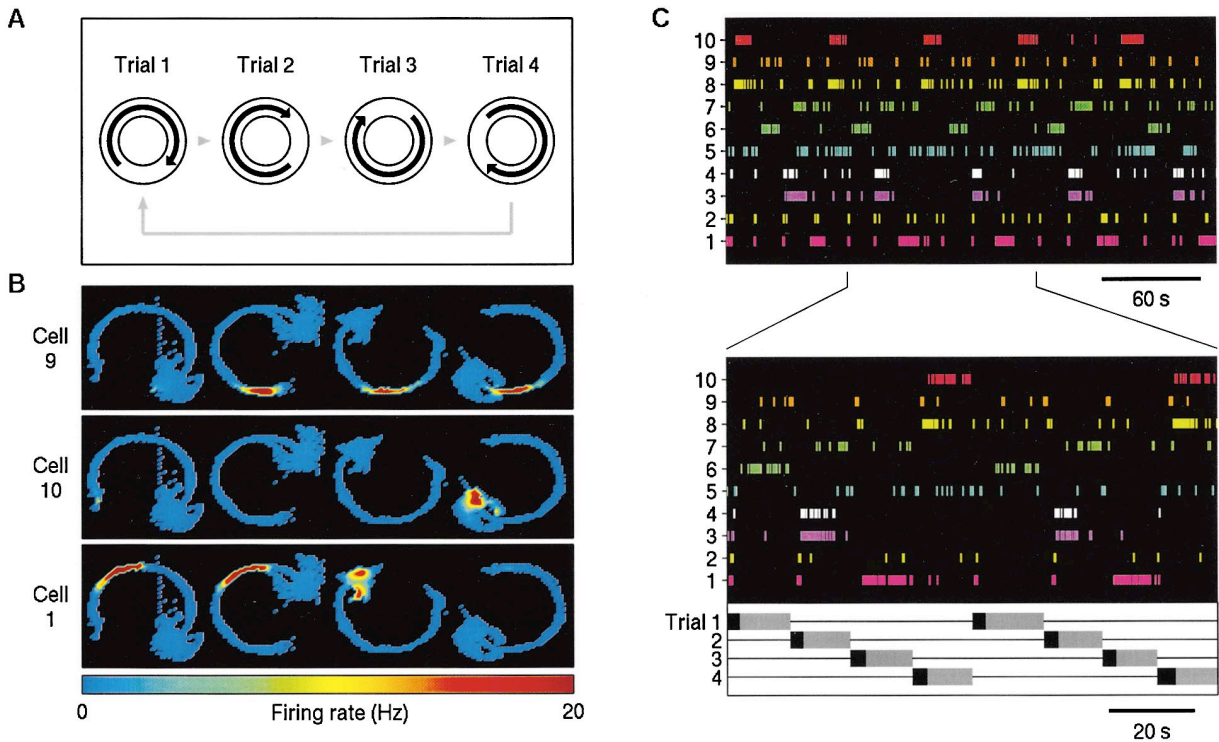


Figure 1. Behavioral Task and Hippocampal Unit Activity

(A) Schematic of the four-trial sequence in the circular track task. A single trial consisted of travel from the start location to a removable food well placed at the goal location, followed by food consumption; in any given trial the goal was located at a position 270° clockwise from the start. After completion of a trial, the goal location became the start location for the subsequent trial. After four trials the animal is at its original starting location and the sequence begins again. A recording session consisted of a sleep epoch (conducted in a separate sleep enclosure), a behavioral epoch (RUN) of 40 trials, and a subsequent sleep epoch.

(B) Spatial firing characteristics of three example CA1 cells. Each column represents activity grouped by trial type.

(C) Periodic repetition of characteristic ensemble spiking pattern. (Top) Ensemble activity over a representative 5 min window of RUN. Each vertical tick mark represents a single action potential. Note the regular repetition of the spatiotemporal pattern that corresponds to a single pass through the four-trial sequence. (Bottom) Expanded segment of RUN epoch ensemble activity. Horizontal bars represent the time course of the four different trial types; black bars denote portions of the trial during which the animal is traversing from start to goal location.

(e.g., cell 10, Figure 1B), similar to behavioral dependence reported in other tasks (Wiener et al., 1989; Deadwyler et al., 1996). Note that the combination of spatial receptive fields and structured spatial behavior produces a characteristic ordered pattern of ensemble activity (Figure 1C). The temporal structure within this pattern is determined by the sequence in which the animal's behavior takes it through the task environment, providing within the ensemble pattern a unique signature of the behavioral experience. Due to the repetitive nature of the task, such patterns of activity were consistently repeated throughout a given session. The repeated activation of these robust patterns during a behaviorally salient task led us to hypothesize that such patterns may be good candidates for subsequent reproduction during sleep.

REM episodes were identified as periods of sleep with sustained (>60 s) increases in the local field potential theta power (quantified in the theta/delta power ratio and confirmed by video monitoring of immobility and sleep posture). The pattern of neuronal ensemble activity over the entire duration of each identified REM window (the template) was then examined for correspondence to patterns recorded during RUN (Figure 2). In

contrast to studies that investigated the recurrence of multineuron spike sequences on the timescale of milliseconds to seconds (Abeles and Gerstein, 1988; Abeles et al., 1993; Nadasdy et al., 1999), we examined neural activity at a lower temporal resolution but over much longer durations on the order of tens of seconds to minutes, with individual neuron spike train data binned at 1 s resolution and Gaussian smoothed ($\sigma = 1.5$ s). This degree of binning and smoothing preserves and emphasizes modulation of neuron activity that occurs at behavioral timescales, such as place field activation, while eliminating millisecond-timescale temporal structure.

To quantify the similarity between a RUN epoch and a given REM episode, we defined a template correlation coefficient (C_t) between two multiple-neuron spiking patterns. If a spatiotemporal pattern of ensemble activity is represented as a matrix with the dimensions of time and cells, C_t between a given REM template and RUN window is analogous to the degree of overlap observed when the two matrices are superimposed. To compare the activity from individual REM episodes to the considerably longer RUN epochs, each REM pattern was used as a sliding template to identify matching patterns within the RUN epoch. As shown in Figure 2B, C_t was calcu-

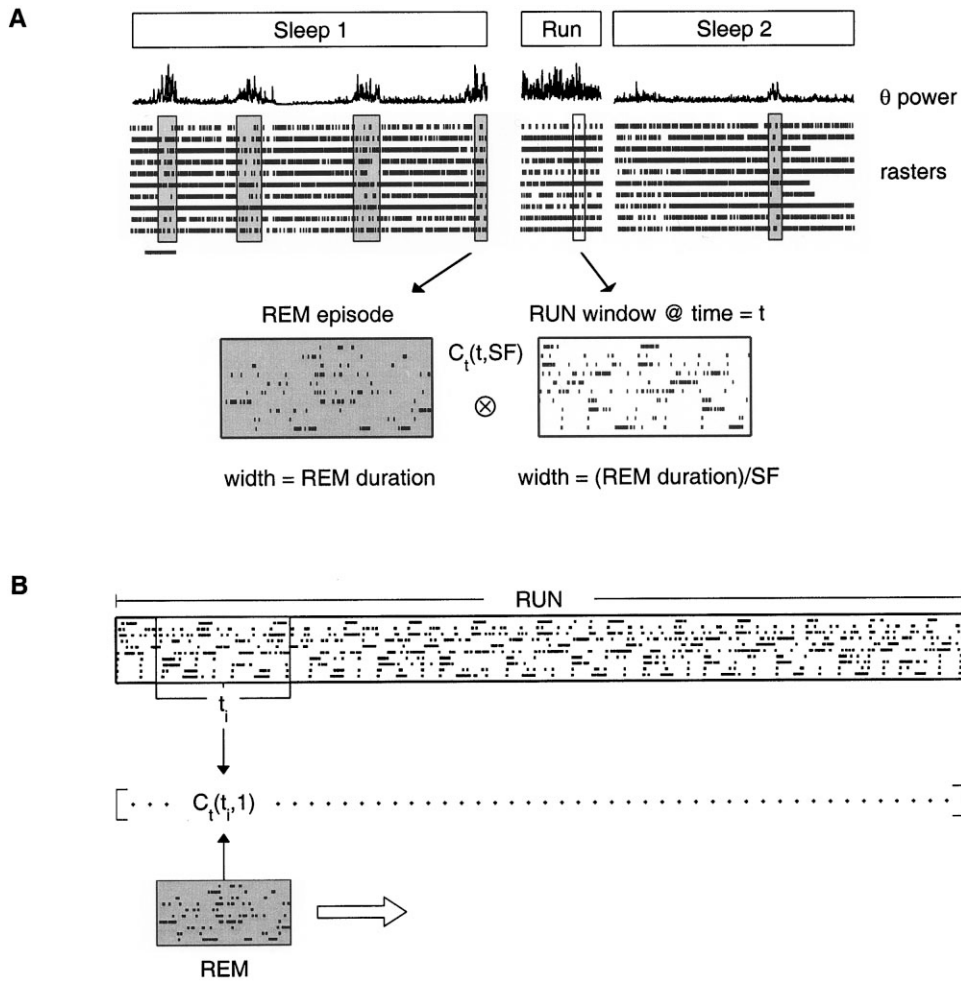


Figure 2. Identification of REM Sleep Templates for Correlation Analysis

(A) Experimental design. REM episodes identified by increases in LFP theta power are used as templates in independent searches across the RUN epoch. The template correlation coefficient (C_t) is calculated between the template and multiple RUN windows in a sliding window fashion. The width of the RUN window is defined by the scaling factor (SF); SF = 1 corresponds to equivalent REM and RUN window lengths, while SF > 1 corresponds to relatively smaller RUN windows (i.e., slower REM activity). Scale bar, 4 min.

(B) Schematic of sliding window correlation analysis. For each time point t_i in the RUN epoch, a window of RUN ensemble activity centered at that time is extracted and compared to the REM template activity. The result is a correlation vector encompassing the entire RUN epoch and signifying the strength of correspondence between the REM template and different points during RUN. Note that temporal scaling is introduced into the correlation by varying the width of the RUN window taken around each time point ($\text{width}_{\text{RUN}} = \text{width}_{\text{REM}}/\text{SF}$). The correlation depicted here represents C_t analysis with SF = 1; correlation was repeated for SFs ranging from 0.3 to 3.0, and the collection of C_t vectors at different SFs defines the C_t matrix.

lated between the REM episode pattern and RUN patterns from windows centered at successive time points across the RUN epoch (step size, 1 s). In addition to identifying proper temporal alignment, evaluation of correspondence requires consideration of temporal scaling because reactivated activity during REM may be compressed or expanded compared to RUN activity. To account for this, the correlation analysis described above was repeated at multiple temporal scaling factors (SF). SFs > 1 signify a slower corresponding activity during REM, while SFs < 1 signify faster activity. The result is a two-dimensional correlation map of the RUN epoch, with each point $C_t(t,SF)$ signifying how strongly a segment of RUN activity centered at time t corresponds to the REM template at a given scaling factor SF. An example of two correlated ensemble patterns is shown in

Figure 3 (120 s REM template and corresponding 75 s RUN window, $C_t = 0.32$).

To establish that observed correlations between REM and RUN patterns could not have arisen by chance alone, the significance of C_t was assessed relative to a sample distribution of shuffled-template correlation data generated for each REM episode. Each REM template was randomized to create a sample of possible templates specific to that REM episode ($n = 50$). The template correlation function C_t was then calculated for every shuffled template to create a distribution of possible C_t values for every (t,SF) point. Shuffles were performed upon binned spike count data prior to Gaussian smoothing. Because no single shuffle procedure is comprehensive, we used four different shuffled C_t functions, each designed to address nonspecific populationwide

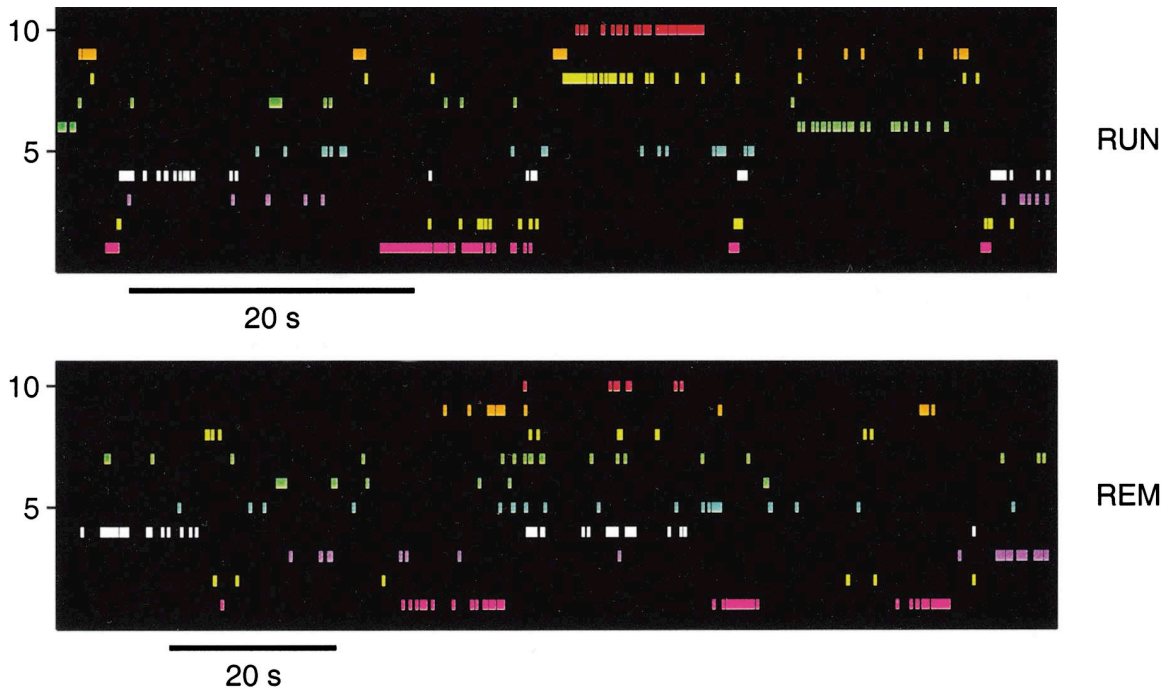


Figure 3. Example Correspondence between a REM Template and RUN Activity
(Top) Rasters of 10 pyramidal cells during a 75 s window from RUN. The RUN time axis is scaled to maximize raster alignment with REM (SF = 1.6). (Bottom) Rasters of the same cells over the duration of a 120 s REM template.

effects that could contribute to measured REM-RUN correlation (Figure 4). First, to control for the possibility that REM-RUN correlation was the result of consistent differences in firing rate between cells, spike count data were independently shuffled within each cell, preserving overall firing rate while disrupting longer timescale temporal structure within and between cells (BIN shuffle). Second, binned spike counts were shuffled in time similar to the BIN shuffle but with relative spike count data across cells held fixed (COLUMN shuffle). This would preserve population vectors that are reactivated as discrete states, like those observed in SWS reactivation (Wilson and McNaughton, 1994), but would disrupt long timescale temporal ordering between states. Nonspecific correlation could also arise due to broad modulation of overall activity in both RUN and REM. Temporally intact spike count vectors were exchanged between cells in the third shuffle (SWAP shuffle). Finally, to ensure that REM-RUN correspondence depended upon the temporal alignment of activity across cells, spike count data for each cell were randomly displaced in time relative to activity in the other cells while maintaining within-cell spike timing information (SHIFT shuffle). This preserves the temporal structure of individual cell firing patterns while disrupting the relative phase between them. It is important to note that significant REM-RUN correlation requires correspondence in both the firing patterns of individual neurons as well as the temporal alignment of activity across all neurons.

At each time point and SF during RUN, the observed correlation function C_t was converted into four z scores relative to the individual shuffled-template distributions. Each C_t value was then associated with an overall z

score, equivalent to the minimum (least significant) z score relative to the four shuffle distributions. Thus the two-dimensional C_t matrix is converted into a two-dimensional z score matrix describing the significance of correspondence between a given REM template and patterns across the RUN epoch (Figure 5A). While it is possible to characterize the significance of individual peaks in the correlation function, we employed a more stringent test incorporating the repetitive structure of the behavioral task itself. Each RUN period was divided into behavioral epochs corresponding to repetitions of the four-trial task sequence, producing a C_t significance matrix for each behavioral segment (Figures 5A and 5B). These matrices were averaged across all behavioral epochs, and the correlation significance of each REM template was defined as the maximum value of the epoch-averaged matrix.

We examined a total of 45 REM episodes from four animals over eight different recording sessions. REM episode durations ranged from 60 s to 250 s (mean, 114.0 ± 50.2 s). There was a notable asymmetry between prebehavior and postbehavior depth of sleep, as quantified by REM episode incidence (prebehavior 3.0 episodes/hr, postbehavior 0.7 episodes/hr) and percentage of time spent in REM (prebehavior 9.3%, postbehavior 2.4%). This difference may be attributable to the animal's behavioral state immediately following task performance. Twenty of 45 (44.4%) REM episodes showed significant correlation to RUN activity ($p < 0.05$, 19/38 prebehavior, 1/7 postbehavior, Figure 5C). Peak correlation significance occurred at temporal scaling factors ranging from 0.55 to 2.49 (mean 1.4 ± 0.6), with the majority (65%) of peak significance points correspond-

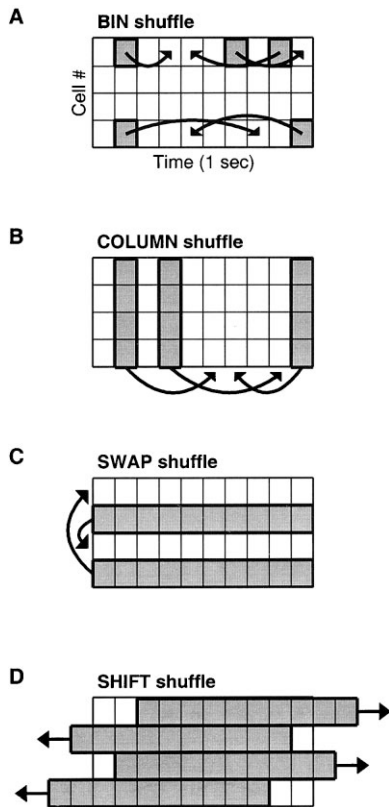


Figure 4. Ensemble Pattern Shuffle Analyses

(A) BIN shuffle. All shuffles performed on binned ensemble spike train data, represented here as a two-dimensional matrix. In the BIN shuffle, bins are pseudorandomly exchanged within each cell spike train vector, with shuffling performed independently on each spike train.

(B) COLUMN shuffle. Similar to the BIN shuffle, except temporal alignment of spike train data is preserved across cells.

(C) SWAP shuffle. Entire spike train vectors are pseudorandomly reassigned between cells. Note that the temporal order of spike activity within each spike train is preserved.

(D) SHIFT shuffle. Entire spike train vectors are temporally shifted relative to original alignment, with relative temporal order preserved within each spike train. The shift distance is pseudorandomly chosen and ranges between half the window length backward and half the window length forward. The shift is circular, such that data removed from the pattern at one end is reinserted at the opposite end.

ing to $SF > 1.0$, suggesting that REM activity recapitulates RUN activity at approximately the same speed or slower.

The reactivation of hippocampal patterns during SWS is strongest immediately following awake behavior, suggesting the development of an experience-dependent memory trace (Wilson and McNaughton, 1994; Kudrimoti et al., 1999). However, we observed significant RUN correlation within 19 of 38 REM episodes recorded before familiar RUN behavior on any given recording day. Does this activity reflect persistent mnemonic activation? Because animals received repeated daily exposure to the task, RUN-correlated patterns during prebehavioral REM sleep may be attributable to residual activity from previous behavioral sessions.

If structured REM activity represents the experience-dependent reactivation of patterns established during

RUN behavior, there should be no significant correlation between REM episodes and novel RUN behaviors to which the animal has never been exposed. We therefore examined three additional experiments where the animal was exposed to both the familiar RUN task as well as a novel spatial task (RUN*, Figure 6A). These novel tasks were also spatial locomotor tasks with multiple food reinforcement points and multiple repetitions (see Experimental Procedures). When we compared prebehavioral REM episodes to these novel RUN epochs, we detected no significant correlation between them (15 REM episodes, Figure 6B). Furthermore, the distributions of correlation significance scores were significantly different in novel versus familiar environments ($p < 0.00005$, Kolmogorov-Smirnov test, Figure 6C). Three REM episodes identified during sleep following RUN* were tested. While none were found to have significant correlation with RUN* epochs, the small number of samples, consistent with earlier observations of limited post-behavioral REM episodes, makes evaluation of this result difficult. This may reflect either a difference in quality of postbehavioral sleep as previously indicated or a slower incorporation of mnemonic information into REM. It is important to note that the same REM episodes that failed to match novel RUN* epochs did exhibit a significant distribution of correlation scores to familiar RUN epochs (identical to the distribution of all REM-familiar RUN correlation scores), demonstrating that the lack of novel RUN* correspondence was not due to a bias in the sample of REM episodes. This suggests that the observed correspondence of REM activity to RUN patterns in the three-quarters circular task arises from the replay of previously learned, behavior-specific activity.

To investigate correspondence between REM sleep and broader characteristics of awake behavior, we next examined variations in the theta rhythm, a large amplitude 6–10 Hz oscillation in hippocampal extracellular field potential regulated by medial septum cholinergic and GABAergic inputs (Vanderwolf, 1969; Stewart and Fox, 1990). The theta rhythm strongly modulates single-cell firing rates and excitability and may be important for the induction of synaptic plasticity. Theta frequency oscillation is prominent during awake behavior and REM sleep and is highly correlated with specific behaviors in different species, such as exploration and movement in rodents (Buzsaki et al., 1983). Because of this behavioral dependence of theta rhythm strength, different segments within a behavioral task will elicit different amounts of theta activity. Hippocampal local field potential (LFP) traces recorded during RUN epochs exhibited phasic increases and decreases in theta rhythm strength that were tightly coupled to the repetition of single trials within the circular track task (Figure 7A).

The strength of the theta oscillation (measured as power of the 6–10 Hz bandpass filtered LFP trace, see Experimental Procedures) was calculated across all REM episodes ($n = 20$) that exhibited significant RUN correspondence in their ensemble unit activity. To examine similarities between REM and RUN patterns of theta frequency modulation, each REM theta power trace was then aligned with its corresponding RUN theta power trace according to the temporal alignment and scaling values determined by template correlation anal-

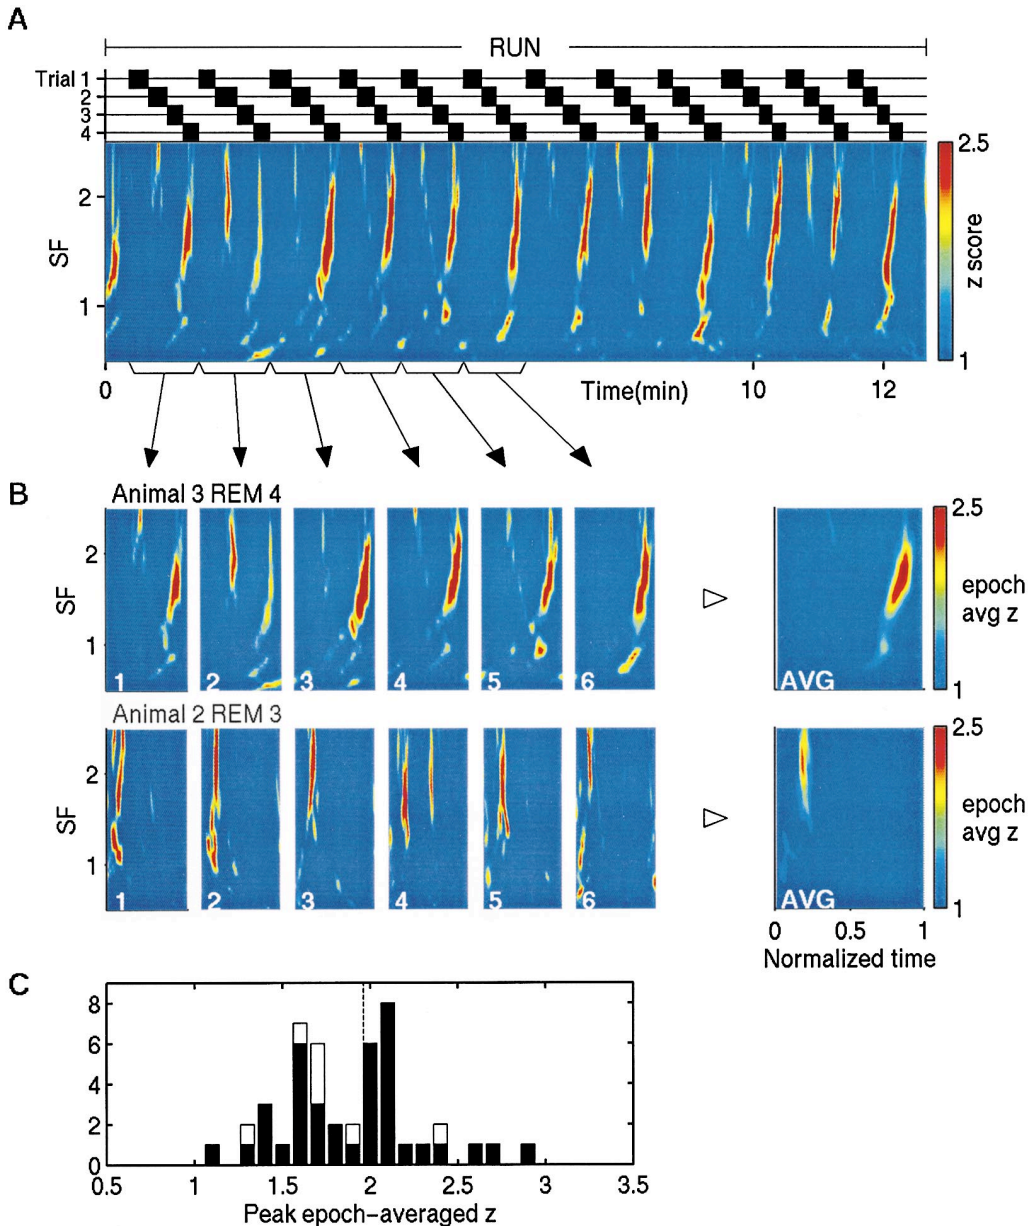


Figure 5. Template Correlation Analysis of REM-RUN Correspondence

(A) Example correlation z score analysis. False-color image represents correlation z score data between one REM episode (animal 3 REM 4) and the entire RUN epoch. The C_i value at each (t,SF) point during RUN is converted to four z scores relative to the shuffled-template distributions; significance of the template correlation coefficient at each point is designated by the minimum z score. The repetition of behavioral trials during RUN are represented in the timeline at top.

(B) Behavioral epoch analysis of two example REM episodes. Each left-hand panel plots the minimum z score across a single repeat of the four-trial behavioral sequence, calculated at temporal scaling factors from 0.5 to 2.5; plots have been normalized along the time axis within each epoch. Data in the top row is from the analysis shown in Figure 4A. Only the first six behavioral epochs are shown. The general correspondence of a REM episode to RUN was evaluated by averaging minimum significance values across repeated behavioral segments. The result is an epoch-averaged z score function for each REM episode, as displayed in the right-most panels.

(C) Distribution of peak epoch-averaged z scores for all REM templates. Black portion of bars, prebehavior REM episodes; white portion of bars, postbehavior REM episodes. Bars to the right of the dashed line denote REM episodes with significant template correlation ($p < 0.05$).

ysis. For example, Figure 7A shows LFP theta power from the correlated REM episode and RUN window depicted in Figure 3. Peaks and troughs were identified in the RUN theta trace (red and blue dots, respectively, Figure 7B); corresponding REM theta values were measured and divided into two groups depending on their

RUN alignment (H, aligns with RUN peak; L, aligns with RUN trough). Mean H and L REM theta power values were calculated for each REM episode and normalized for comparison across all REM episodes. In 75% of REM episodes (16/20), the mean H theta power value was greater than the mean L theta value; furthermore, mean

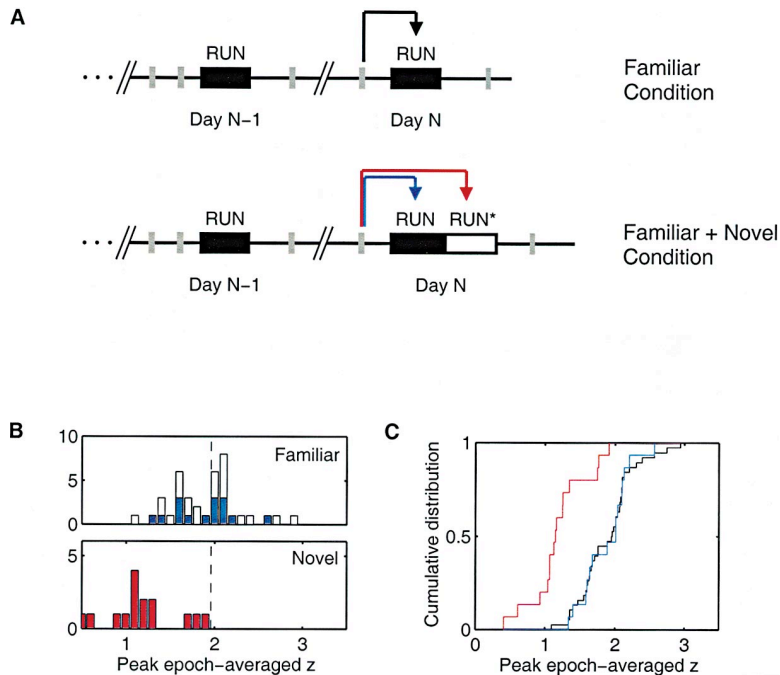


Figure 6. REM Correspondence to Novel Versus Familiar RUN Epochs

(A) Schematic of recording session time course. Black bars, familiar RUN epochs. White bar, novel RUN* epochs. Gray bars, REM episodes. Recording sessions are separated by ~24 hr, as indicated by the diagonal lines. Note that REM episodes occurring before a familiar RUN epoch on any given day actually follow the previous day's RUN epoch. To investigate the experience dependence of prebehavior REM correspondence to familiar RUN behaviors, ensemble activity from REM episodes were also compared to neural patterns recorded during novel behaviors (RUN*). The extent of this novel RUN*-REM correspondence (red arrow) can be compared to familiar REM-RUN correspondence, both for the same REM episodes used in the novel analysis (blue arrow) and for all other prebehavior REM episodes (black arrow).

(B) Distributions of epoch-averaged correlation z scores from REM episodes recorded before novel RUN* versus familiar RUN behaviors. Scores greater than the dashed vertical line indicate a significant correlation ($p < 0.05$). (Top) Correlation scores between REM episodes and familiar RUN epochs. White portion of bars, all prebehavior REM episodes during novel condition; blue portion of bars, prebehavior REM episodes during familiar and novel condition. (Bottom) Correlation scores between REM episodes and novel RUN* epochs.

(C) Cumulative distributions of correlation significance scores. Red line, significance of correlation to novel RUN* behaviors. Blue line, significance of correlation to familiar RUN behaviors, same REM episodes as novel data. Black line, significance of correlation to familiar RUN behaviors, all prebehavior REM episodes. The distribution of correspondence to novel RUN* behaviors differs significantly from the distributions of correspondence to familiar RUN behaviors, both for the subset of REM episodes tested against novel behaviors as well as for all REM episodes ($p < 0.00005$, Kolmogorov-Smirnov).

H and L theta power values averaged across all evaluated REM episodes are significantly different ($p < 0.0005$, paired t test, Figure 7C). This significant difference between REM theta values that were divided according to their alignment with RUN theta values suggests that aspects of theta oscillation modulation generated during the awake behavioral task are also represented during REM sleep.

Discussion

The gradual shift in the locus of memory storage from the hippocampus to other, presumably neocortical, sites suggests that previously stored memories can be subsequently reactivated (Squire, 1992). Here, we demonstrate significant ensemble correlation between periods of awake behavior and REM sleep, despite the absence in REM of the explicit sensorimotor cues that drive distinct neural patterns during RUN. The existence of decipherable mnemonic structure during REM sleep raises further questions regarding the neural mechanisms responsible for such temporally structured activity, as well as the possible role of such reactivation in processes such as memory consolidation and learning.

Specificity of REM-RUN Correspondence

Analysis of REM-RUN correspondence demonstrated that the temporal patterns of individual neuronal spiking and the phase or timing of firing between different neurons established during RUN are recapitulated during

REM. This analysis employed a template correlation measure (C_t) that quantified the strength of similarity between two patterns of activity. The crucial question that must be asked is whether REM-RUN correspondence is a specific result of behavioral experience or whether such similarity could arise due to nonspecific patterns of activity. In this paper we have addressed this issue through the use of shuffled template variants and the examination of REM correspondence with novel patterns of RUN activity.

Shuffle procedures were selected to control for several potential nonspecific sources of correspondence (Figure 4). The BIN shuffle addresses the possibility that correspondence could have arisen from general equivalence of individual cell firing rates between REM and RUN. Because both states are marked by increases in theta rhythmicity, cells that systematically changed firing rate during theta modulated states could contribute to REM-RUN correspondence. This shuffle preserves relative firing rates between cells but disrupts the temporal patterns that are a direct consequence of the interaction between behavior and place specificity of firing. This type of nonspecific rate effect was also controlled for by the novel RUN* analyses, since both the familiar and novel behaviors generated prominent theta rhythmic activity, but significant correspondence occurred only in the familiar condition.

A potential source of nonspecific match between RUN and REM is the presence of discrete episodes of characteristic activity that were not related to the specific RUN

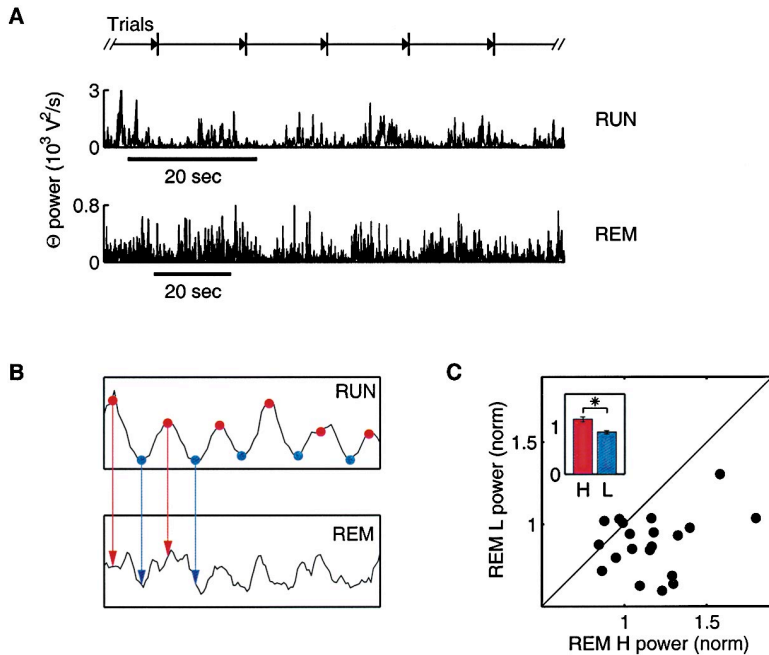


Figure 7. REM-RUN Correspondence in Theta Rhythm Modulation

(A) Broad patterns of modulation in the theta rhythm. (Top trace) LFP theta frequency power during the 80 s RUN window displayed in Figure 3, with theta power evaluated as the squared amplitude of the 6–10 Hz bandpass filtered LFP signal. Line above the RUN trace denotes the starting and ending points of individual behavioral trials; note the regular phasic modulation of RUN theta power by behavior. (Bottom trace) LFP theta frequency power during the 120 s REM episode displayed in Figure 3. Theta power traces are aligned and scaled based on template correlation analysis, i.e., at the time and scaling factor corresponding to maximal ensemble pattern correlation derived from unit rasters (maximum C_i).

(B) Evaluation of REM-RUN theta power correspondence. REM and RUN theta power data are binned at 1 s intervals are aligned and scaled to the optimal values derived from template correlation analysis (Figure 4). REM theta power values that aligned to either peaks or troughs in the RUN theta trace (red and blue dots, respectively) were identified and grouped according to their alignment

with RUN (H, REM theta values that aligned with RUN theta peaks; L, REM theta values that aligned with RUN theta troughs). Example H and L values depicted by red and blue arrows, respectively.

(C) REM-RUN theta power correspondence in all REM episodes with significant ensemble correspondence. Each point plots the H theta power versus the L theta power for a single REM episode; values are normalized within REM episodes by the mean theta power amplitude for comparison across episodes. Under the null hypothesis of independence between REM and RUN theta power modulation, there should be no difference between H and L REM theta values calculated from template correlation-derived alignment to RUN data. (Inset) Mean H and L theta power (\pm SEM) across all REM episodes. There is a significant difference between REM H and L values ($p < 0.0005$, paired t test), suggesting that aspects of theta oscillation modulation generated during the awake behavioral task are represented during REM sleep.

experience. For example, occasional bursts in synchronized activity across the hippocampus can occur due to normal (large irregular activity) phenomena; the emergence of such phenomena in both RUN and REM could lead to apparent correspondence driven by synchronous population activity. Several of the shuffling procedures directly address this possibility. The SWAP shuffle exchanges the identity of individual cells but maintains the time course of any populationwide modulation that might exist. If populationwide covariance in activity is the source of apparent correspondence, the precise identity of the cell is less significant than the proper temporal alignment of all cells in the ensemble during these discrete events. The COLUMN shuffle preserves the ensemble structure of activity within discrete windows but alters the temporal order of these windows. Correspondence resulting from the appearance of discrete events would be preserved in this shuffle while patterns that are dependent upon the temporal ordering of events across windows, such as patterns of place-related firing, would be disrupted.

Slow rhythmic modulation of population neural activity, such as the cortical slow oscillation, is known to occur during sleep. Another possible nonspecific explanation for the observed match between RUN and REM patterns is that such broad fluctuations in overall neural activity during RUN could match nonspecific slow rhythmicities expressed during REM. However, such slow fluctuations are not likely to account for the observed correspondence between specific ensemble spiking

patterns for several reasons. First and foremost, the potential confound introduced by an underlying slow rhythm is periodic activation of entire subpopulations of cells, a result that is specifically controlled for within the shuffle analysis by the COLUMN and SWAP analyses. Second, place specificity of firing of individual cells clearly demonstrates that neurons were not firing in a way that simply reflected behavioral state. Place-specific cells fired at specific points within each trial, and cells with different place fields fired at different times; accordingly, RUN ensemble patterns are composed of neuronal activity with specific temporal offsets rather than phasic activation of the entire ensemble (Figure 1). Template correspondence depends then on not just periodic activation within the REM episode but periodic activation with correct temporal offsets, a result that is not easily explained by an underlying slow rhythm. While existence of slow rhythms that also imposed consistent phase relationships between the firing patterns of different neurons during both RUN and REM cannot be ruled out, such activity has never been demonstrated and would be extremely difficult to reconcile with the observed spatial specificity (place fields) of individual neurons.

The examination of novel RUN* behaviors further demonstrates the dependence of REM-RUN correspondence on the specific patterns of ensemble activity generated during experience in the familiar testing environment. We examined several tasks where similar periodicities in behavior were produced due to the repetitive

nature of the task. To explicitly control for the possibility that changes in apparatus might produce subtle alterations in behavior that would impair REM match, we examined a control task in which the animal simply altered the pattern of starting and stopping locations while still remaining on the familiar apparatus (RUN*3). This preserves the overall quality of behavior (circular, periodic running) as well as the general periodicity (~15 s per trial) but significantly alters the precise pattern of neural activity by proceeding through a different sequence of locations within the apparatus. The clear difference in the degree of correspondence between REM patterns and any and all of these novel conditions (Figure 6) indicates that ensemble match between familiar RUN and REM was not a simple consequence of non-specific regularities in neural activity in the familiar RUN due to such factors as periodic fluctuations in behavior.

The robust spatial correlates of hippocampal neuronal activity (place fields) indicate that the activity of these cells is a function of spatial location, but the apparent periodicity of firing across the RUN task raises the possibility that these cells have simply been entrained into firing with a slow periodicity that leads to the appearance of spatially specific firing. If the consequence of RUN activity was to reinforce slow periodic firing patterns of individual cells and the tendency to fire with similar slow periodicity was reflected during REM, it could be argued that C_t correspondence was not due to a match with the behaviorally specific ensemble pattern of RUN-related firing but simply the periodicity of firing of individual cells. This possibility was addressed through the use of the SHIFT shuffle. The signature ensemble activity that characterizes a specific RUN episode is dependent not only upon the temporal pattern of firing of individual cells, but also upon the relative timing or phase of firing between cells. The SHIFT shuffle preserves the temporal firing patterns of individual cells but disrupts the phase information between cells. The demonstration that significant REM-RUN correspondence is only achieved when both the firing patterns of individual cells as well as their phase relationship with other cells is preserved indicates that this correspondence is due to the explicit match with the behaviorally specific patterns of ensemble activity produced during RUN.

In contrast to experience-dependent changes observed during SWS, significant RUN-correlated activity occurred during REM episodes prior to awake behavior. Does this correspondence reflect previously encoded memories? Examination of neural activity during REM episodes preceding three different novel behaviors revealed no RUN correspondence, despite the fact that approximately half of the same REM episodes were significantly correlated to patterns from the familiar task. Thus, RUN-correlated REM patterns recorded before task performance on any given day represent persistent experience-dependent activity from previous task sessions. This time course differs from that of mnemonic changes in SWS, which are strongest immediately following awake behavior. However, some residual SWS reactivation can be observed in prebehavioral sleep (Kudrimoti et al., 1999), and behavioral studies of experience-dependent changes involving REM sleep mirror this longer time course. Increases in REM following learning occur as much as 24 hr after the end of training,

and acquisition of certain memory paradigms requires REM sleep hours to days after learning (Smith, 1995). The apparent difference in the robustness of reactivation between pre- and postbehavioral REM episodes may either reflect a simple difference in the quality of REM between these periods that consequently limits the expression of reactivation or may reflect differences in the processing of mnemonic information within the sleep cycle. While present data suggests a difference in REM reactivation during these periods, further study will be required to evaluate the significance of this effect.

Neural Mechanisms

These results demonstrate that the relative temporal firing order within an assembly of neurons can be preserved and reproduced. It is critical to note that the timescale of these temporal patterns extended over tens of seconds to minutes. In contrast to previous studies of temporal sequence activity that examined sequences on the timescale of milliseconds to seconds (Abeles and Gerstein, 1988; Abeles et al., 1993; Nadasdy et al., 1999), reactivation of temporal sequences at this timescale has never been previously observed. Previous studies that failed to identify behaviorally related ensemble activity during REM looked at short latency correlation but did not examine long timescale temporal structure (Kudrimoti et al., 1999). The surprising length of reactivated sequences raises the question of how temporal information at such a scale is encoded. Hippocampal CA1 place fields develop a strong asymmetry with experience, which provides a synaptic mechanism capable of encoding a sequence of locations (Mehta et al., 2000). Sequence information is crucial for generating neural activity dependent upon temporal order, such as the trajectory-dependent cell firing observed in the hippocampus and entorhinal cortex (Frank et al., 2000; Wood et al., 2000). Trajectories represent task-specific temporally ordered spatial locations, and linkage of such trajectory representations either within the hippocampus or entorhinal cortex could provide a mechanism for reconstructing extended sequences of behavior.

The establishment of temporal order may involve extrahippocampal brain areas such as neocortex. In particular, the prefrontal cortex can play a role in maintaining information relating temporally adjacent states such as the beginning and end of a trial. Neurons in the prefrontal cortex have broad behavioral correlates during spatial behavioral tasks (Jung et al., 1998) and exhibit prospective activity that can encode temporal relationships across delay periods (Quintana and Fuster, 1992; Watanabe, 1996; Asaad et al., 1998; Rainer et al., 1999). Coordinated interactions between the hippocampus and prefrontal areas during REM sleep could provide a mechanism for organizing temporal order of hippocampal or entorhinal states representing behavioral sequences.

Given the subcortical generation of the theta rhythm, the recapitulation of patterns of theta modulation in addition to ensemble patterns of pyramidal cell activity suggests a broad recapitulation of behavioral state. Interestingly, in accordance with the mild temporal expansion of REM reactivation suggested by the occurrence of optimal scaling factors between 1 and 2, the fre-

quency of theta during REM sleep is ~ 1.2 times slower than during RUN (data not shown). This approximate temporal concurrence between theta rhythm and REM-RUN correspondence may reflect globally slower neural processing during sleep. For example, the frequency of the theta rhythm is sensitive to brain temperature (Whishaw and Vanderwolf, 1971) and brain temperature is typically lower during sleep (Andersen and Moser, 1995), suggesting that the neural processes underlying REM reactivation may be similarly slowed. Alternatively, there may exist a specific link between the theta rhythm and sequence reactivation, with the theta rhythm serving as a pacing mechanism during temporal information storage and reactivation, perhaps to coordinate interactions across multiple brain regions. Further experiments will be required to explore the role of the theta rhythm in temporal scaling.

Functional Implications

What could be the function of REM sleep replay of awake activity? One possible interpretation is that REM activity reflects neocortical activation of hippocampal circuits in a later stage of the memory consolidation process (Hennevin et al., 1995; Stickgold et al., 2000). The reactivation of short-timescale hippocampal patterns is enhanced during periods of SWS immediately following behavior (Wilson and McNaughton, 1994; Kudrimoti et al., 1999). In particular, the synchronization of subsets of hippocampal neurons during oscillatory ripple events has been suggested as a strong mechanism for synaptic modification (Buzsaki, 1989). Recently acquired information within the hippocampus may activate neocortical circuits as the initial stage of consolidation, as suggested by the correlation between neocortical spindle activity and high-frequency hippocampal discharges during SWS (Siapas and Wilson, 1998). Neocortical circuits established via SWS hippocampal-neocortical interactions could subsequently engage the hippocampus during REM sleep in the form of long timescale sequences that last for minutes. This reactivation of previous behavioral episode representations may be important for the learning and performance of procedural tasks, which is dependent upon REM sleep (Karni et al., 1994). Mnemonic information that may have shared characteristics along a particular behavioral axis such as emotion could be juxtaposed and evaluated for common causal links, allowing adaptive behavioral change based on prior experience (Hobson et al., 1998). The ability to identify specific mnemonic content within REM sleep will allow explicit evaluation of such hypotheses and further the examination of the role of sleep and dreaming in memory formation and consolidation.

Experimental Procedures

Behavioral Paradigm

Following implantation, Long-Evans rats (5–7 months old) were trained to run from a start location to a goal location for a food reward on an elevated circular track (95 cm diameter, 10 cm width). This task (RUN) was considered familiar because all animals were trained daily on the task for at least 5 days prior to the first recording session. A trial consisted of travel from the start location to a removable food well placed at the goal location, followed by food consumption; in any given trial the goal was located at a position 270° (clockwise) from the start. After completion of a trial the goal location

became the start location for the subsequent trial. After four trials the original start location once again became a start position, and the entire four-trial sequence was repeated. Note that because the animal was always traversing toward a location with a food well, no explicit behavioral criterion for task performance was necessary other than steady, consistent locomotion without interruption between start and goal locations. A recording session consisted of a 1–2 hr sleep epoch, a 10–15 min behavioral epoch (RUN) of ~ 40 trials, and a subsequent sleep epoch. All sleep sessions were conducted in a separate sleep enclosure within the recording room. Recording sessions were conducted daily, always at the same time of day such that ~ 18 – 20 hr separated the end of one session with the beginning of the subsequent session.

Three additional novel spatial locomotor tasks (RUN*) were conducted in the fourth experimental animal to directly address the issue of experience-dependence. Novel RUN* behaviors were chosen to be similar to the familiar RUN task, and each consisted of a spatial locomotion task between food reward sites with trials repeated throughout the RUN epoch. RUN*1 was performed on an elevated T-shaped track (115 cm central arm, 60 cm choice arms), and a behavioral trial consisted of food travel from the goal location on the central arm to a goal location on one of the arms, food consumption, and then return to the central goal location. The animal tended to perform this task in a delayed alternation pattern, alternating between left and right arms in subsequent trials, and analysis was restricted to behavioral segments where the animal reliably visited one arm followed by the other arm. RUN*2 was performed on an elevated U-shaped track (160 cm arms, 15 cm connector), and a behavioral trial consisted of travel from one goal location down one arm, across, and up the other arm to the opposite goal location, food consumption, and return to the original goal location. RUN*3 was performed on the same circular track as the familiar RUN task, and the behavioral task was identical to RUN except the animal traveled 90° during each behavioral trial.

Electrophysiology

Following surgical implantation with a microdrive array of 12 independently adjustable tetrode wires (AP -3.6 , L 2.2), tetrodes were lowered to the CA1 layer over a period of days and individually positioned to obtain maximal unit isolation. Electrical signals were passed through two miniature 25-channel head-stage preamplifiers to 8-channel differential amplifiers (Neuralynx), bandpass filtered (300 Hz to 6 kHz), sampled (31.25 kHz/channel), and digitized; suprathreshold events were stored for subsequent analysis. Continuous LFP recordings were obtained from a subset of the tetrodes used for unit recording (filtered at 0.1 Hz to 475 Hz, sampled at 1.5 kHz/channel). Head position and direction during RUN epochs were monitored at 30 Hz with a spatial resolution of 0.5 cm via overhead camera tracking of a head-stage infrared diode array. A custom software package (Xclust, M. A. W.) was used to identify clusters of spike waveforms using spike width and peak amplitude on each of the four tetrode channels as primary waveform parameters.

Putative pyramidal cells and interneurons were differentiated based on bursting (complex spiking) and waveform characteristics. To restrict our analysis to patterns of RUN activity composed of actively firing units, only pyramidal cells with mean RUN epoch firing rates above 0.2 Hz were included for subsequent analysis. Cells that were not cleanly isolated or with unstable waveforms over the 4–6 hr recording period were excluded. The number of cells that met these criteria in the eight recording sessions ranged between 8 and 13 per session. While the number of cells in the ensembles used in the present analysis is somewhat lower than those found in earlier studies using this technique (Wilson and McNaughton, 1993), it should be noted that the present study placed additional constraint on unambiguous isolation over the extended recording session and restricted sampling to cells active during RUN periods, which typically make up $\sim 30\%$ of available cells. No other selection bias that might have influenced temporal sequence expression or detection was used to identify cells.

To identify REM episodes, LFP traces were digitally bandpass filtered in the delta (2–4 Hz) and theta (6–10 Hz) bands, and power in each band was computed as the time-averaged squared amplitude of the filtered trace. REM episodes were identified as periods

of elevated theta-delta power ratio (> 2.0). To examine correspondence between long duration patterns and to reduce the detection of false-positive correlations associated with short duration patterns, we limited our analysis to REM episodes longer than 60 s in duration. Sleep during these intervals was verified on videotape.

Template Correlation Analysis

For a given REM template, the spike count data recorded simultaneously from C cells can be transformed into a $C \times N$ matrix, where N is the length of the window in time bins (1 s). Ensemble pattern correspondence can then be evaluated by examining the overlap between the $C \times N$ REM matrix and a corresponding $C \times N$ matrix of RUN activity.

To create the REM template array, the REM episode spike train of each unit was binned into 1 s intervals and the resultant vector was smoothed (convolution with Gaussian, $\sigma = 1.5$ s). The multicell collection of binned and smoothed spike train vectors defined the REM firing rate array. For a given time point t in RUN, the analogous RUN array was constructed over a temporal window centered on that time point; each unit spike train was binned into the same number of intervals as the REM array and the vectors were Gaussian smoothed. The template correlation coefficient (C_t) is defined as:

$$C_t = \frac{1}{N \cdot C} \sum_{c=1}^C \sum_{n=1}^N \frac{(X_{nc} - \bar{x})(Y_{tc} - \bar{y})}{\sigma_x \sigma_y}$$

where: $X_{1c}, X_{2c}, \dots, X_{Nc}$ and $Y_{1c}, Y_{2c}, \dots, Y_{Nc}$ are the binned smoothed spike counts for cell c in the REM and RUN windows, respectively; N is the number of bins; and C is the number of cells. Each cell vector is normalized by its respective dimensionless root mean square amplitude X_c or Y_c ; \bar{x} and \bar{y} are the mean bin values in the REM and RUN arrays, and σ_x and σ_y are the REM and RUN standard deviations of the normalized binned spike counts across all cells, where:

$$X_c = \sqrt{\frac{1}{N} \sum_{n=1}^N X_{nc}^2} \cdot \text{spikes}^{-1}$$

$$\bar{x} = \frac{1}{N \cdot C} \sum_{c=1}^C \sum_{n=1}^N \frac{X_{nc}}{X_c}$$

$$\sigma_x = \sqrt{\frac{1}{N \cdot C} \sum_{c=1}^C \sum_{n=1}^N \left(\frac{X_{nc}}{X_c} - \bar{x} \right)^2}$$

Note that C_t defined in this manner ranges between -1 and 1 .

A high C_t value at time t indicates that the REM template strongly matches an equivalent RUN window centered at time t . C_t was then evaluated between the REM template and multiple RUN windows across the behavioral session in a sliding window fashion (step size = 1 s). To account for temporal expansion or compression of replayed activity in REM, the fixed length REM template was compared to varying length RUN windows at each time point. The size of the RUN window was determined as the length of the REM episode divided by a temporal SF, which ranged from 0.3 to 3.0. Defined in this manner, SF values > 1 denote temporal expansion of activity and slower replay during REM as compared to RUN; SF values < 1 denote temporal compression and faster replay. Thus, template correlation analysis produces a two-dimensional map of RUN activity representing the strength of a given REM template's correlation to RUN at any given time point and scaling factor.

Theta Rhythm Modulation Analysis

To quantify the time-varying strength of theta frequency oscillation, REM and RUN theta power traces were computed at a 1 s resolution as described above from the LFP signal that exhibited the strongest theta oscillation. REM episode theta power traces were individually aligned with corresponding RUN traces, as dictated by the temporal alignment and scaling values indicated by the template-correlation identified peak epoch-averaged z score. Peaks and troughs were identified in the RUN theta power trace, and REM theta power values that aligned with either peaks and troughs with a ± 1.5 s accuracy were identified and separated into two groups designated "H" and "L" based on RUN alignment. H REM theta values aligned with peaks in the RUN theta power trace, while L REM theta aligned with troughs. Thus, a mean H and mean L theta power value was

calculated for each REM episode. To allow comparison across REM episodes, REM theta power values were normalized by the mean theta power in each individual REM episode.

Acknowledgments

We thank W.F. Asaad, J.A. Dani, A. Lee, G. Liu, E. Lubenov, M.R. Mehta, E.K. Miller, S. Seung, and A.G. Siapas for helpful discussions.

Received October 10, 2000; revised December 5, 2000.

References

- Abeles, M., Bergman, H., Margalit, E., and Vaadia, E. (1993). Spatio-temporal firing patterns in the frontal-cortex of behaving monkeys. *J. Neurophysiol.* **70**, 1629–1638.
- Abeles, M., and Gerstein, G.L. (1988). Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *J. Neurophysiol.* **60**, 909–924.
- Andersen, P., and Moser, E.I. (1995). Brain temperature and hippocampal function. *Hippocampus* **5**, 491–498.
- Asaad, W.F., Rainer, G., and Miller, E.K. (1998). Neural activity in the primate prefrontal cortex during associative learning. *Neuron* **21**, 1399–1407.
- Bliss, T.V.P., and Collingridge, G.L. (1993). A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* **361**, 31–39.
- Buzsaki, G. (1989). Two-stage model of memory trace formation: a role for "noisy" brain states. *Neuroscience* **31**, 551–570.
- Buzsaki, G., Leung, L.W., and Vanderwolf, C.H. (1983). Cellular bases of hippocampal EEG in the behaving rat. *Brain Res.* **287**, 139–171.
- Crick, F., and Mitchison, G. (1983). The function of REM sleep. *Nature* **304**, 111–114.
- Deadwyler, S.A., Bunn, T., and Hampson, R.E. (1996). Hippocampal ensemble activity during spatial delayed-nonmatch-to-sample performance in rats. *J. Neurosci.* **16**, 354–372.
- Frank, L.M., Brown, E.N., and Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* **27**, 169–178.
- Hennevin, E., Hars, B., Maho, C., and Bloch, V. (1995). Processing of learned information in paradoxical sleep: relevance for memory. *Behav. Brain Res.* **69**, 125–135.
- Hobson, J.A., Pace-Schott, E.F., Stickgold, R., and Kahn, D. (1998). To dream or not to dream? Relevant data from new neuroimaging and electrophysiological studies. *Curr. Opin. Neurobiol.* **8**, 239–244.
- Jung, M.W., Qin, Y., McNaughton, B.L., and Barnes, C.A. (1998). Firing characteristics of deep layer neurons in the prefrontal cortex in rats performing spatial working memory tasks. *Cereb. Cortex* **8**, 437–450.
- Kami, A., Tanne, D., Rubenstein, B.S., Askenasy, J.J., and Sagi, D. (1994). Dependence on REM sleep of overnight improvement of a perceptual skill. *Science* **265**, 679–682.
- 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.
- Mehta, M.R., Quirk, M.C., and Wilson, M.A. (2000). Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* **25**, 707–715.
- Nadasdy, Z., Hirase, H., Czurko, A., Csicsvari, J., and Buzsaki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *J. Neurosci.* **19**, 9497–9507.
- O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**, 171–175.
- Pavlidis, 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.
- Poe, G.R., Nitz, D.A., McNaughton, B.L., and Barnes, C.A. (2000). Experience-dependent phase-reversal of hippocampal neuron firing during REM sleep. *Brain Res.* **855**, 176–180.

- Quintana, J., and Fuster, J.M. (1992). Mnemonic and predictive functions of cortical neurons in a memory task. *Neuroreport* 3, 721–724.
- Rainer, G., Rao, S.C., and Miller, E.K. (1999). Prospective coding for objects in primate prefrontal cortex. *J. Neurosci.* 19, 5493–5505.
- Siapas, A.G., and Wilson, M.A. (1998). Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron* 21, 1123–1128.
- Skaggs, W.E., and McNaughton, B.L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* 271, 1870–1873.
- Smith, C. (1995). Sleep states and memory processes. *Behav. Brain Res.* 69, 137–145.
- Squire, L.R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* 68, 185–231.
- Stewart, M., and Fox, S.E. (1990). Do septal neurons pace the hippocampal theta rhythm? *Trends Neurosci.* 13, 163–168.
- Stickgold, R., Whidbee, D., Schirmer, B., Patel, V., and Hobson, J.A. (2000). Visual discrimination task improvement: A multi-step process occurring during sleep. *J. Cogn. Neurosci.* 12, 246–254.
- Vanderwolf, C.H. (1969). Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalogr. Clin. Neurophysiol.* 26, 407–418.
- Watanabe, M. (1996). Reward expectancy in primate prefrontal neurons. *Nature* 382, 629–632.
- Whishaw, I.Q., and Vanderwolf, C.H. (1971). Hippocampal EEG and behavior: effects of variation in body temperature and relation of EEG to vibrissae movement, swimming and shivering. *Physiol. Behav.* 6, 391–397.
- Wiener, S.I., Paul, C.A., and Eichenbaum, H. (1989). Spatial and behavioral correlates of hippocampal neuronal activity. *J. Neurosci.* 9, 2737–2763.
- Wilson, M.A., and McNaughton, B.L. (1993). Dynamics of the hippocampal ensemble code for space. *Science* 261, 1055–1058.
- Wilson, M.A., and McNaughton, B.L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science* 265, 676–679.
- Wood, E.R., Dudchenko, P.A., Robitsek, R.J., and Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* 27, 623–633.
- Zola-Morgan, S., and Squire, L.R. (1993). Neuroanatomy of memory. *Annu. Rev. Neurosci.* 16, 547–563.