

Memory of Sequential Experience in the Hippocampus during Slow Wave Sleep

Albert K. Lee¹ and Matthew A. Wilson¹
Picower Center for Learning and Memory
RIKEN-MIT Neuroscience Research Center
Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Summary

Rats repeatedly ran through a sequence of spatial receptive fields of hippocampal CA1 place cells in a fixed temporal order. A novel combinatorial decoding method reveals that these neurons repeatedly fired in precisely this order in long sequences involving four or more cells during slow wave sleep (SWS) immediately following, but not preceding, the experience. The SWS sequences occurred intermittently in brief (~100 ms) bursts, each compressing the behavioral sequence in time by approximately 20-fold. This rapid encoding of sequential experience is consistent with evidence that the hippocampus is crucial for spatial learning in rodents and the formation of long-term memories of events in time in humans.

Introduction

Information processing in the brain is believed to require coordinated activity across many neurons. However, the exact nature of this code is still largely unknown and remains a fundamental problem in neuroscience. With the recent development of methods to simultaneously record the spiking activity of large numbers of individual neurons, the search for complex firing patterns across multiple cells that could directly address this issue has become possible (Abeles and Gerstein, 1988; Aertsen et al., 1989; Abeles, 1991; Meister et al., 1991; Abeles et al., 1993; Wilson and McNaughton, 1993, 1994; Skaggs and McNaughton, 1996; Kudrimoti et al., 1999; Nádasdy et al., 1999; Wessberg et al., 2000; Louie and Wilson, 2001). Here, we present a novel method to analyze sequential firing patterns involving an arbitrary number of neurons and use it to decode the activity of a population of rat hippocampal neurons during slow wave sleep (SWS).

The hippocampus has been a target of particular interest due to its involvement in the formation of long-term memories of events in time in humans (Scoville and Milner, 1957; Milner, 1966; Zola-Morgan et al., 1986) and spatial memory (O'Keefe and Nadel, 1978; Morris et al., 1982) as well as more general sequence memory (Kesner et al., 2002; Fortin et al., 2002) in rodents. It sits atop a neuroanatomical hierarchy, receiving highly processed information from many areas of the brain, and is thus in an ideal position to link events across space and

time (Mishkin et al., 1998). A rodent hippocampal CA1 pyramidal cell (place cell) fires selectively when the animal is in a particular location (the cell's place field) of an environment (O'Keefe and Dostrovsky, 1971). Since different place cells fire selectively in different locations, a rat's experience in moving from one location to another can be represented by the resulting sequence of place fields traversed. Such a sequence may also be a good model for the temporal order of events experienced by humans. We tested whether repetitive experience of a particular spatial sequence resulted in detectable memory traces of that sequence during SWS immediately afterwards.

A major reason for analyzing neural activity during sleep is that, because of the clear absence of ongoing spatial behavior, the effect of previous spatial experience may be more easily identified. Specifically, by comparing the activity in sleep before and after a task, one may identify experience-dependent changes that could correspond to learning. Here, we find recurring sequences during SWS that involve four or more cells and match the rat's immediately preceding spatial experience. In contrast, there is no indication of any matching sequential structure in SWS immediately before the experience. Together with additional analyses, this provides direct neural evidence of the rapid encoding of extended spatial sequences.

Results

Adult rats ($n = 3$) ran back and forth (23–35 round trips in 20–25 min) along linear tracks (180 cm straight or 450 cm U shaped) for chocolate reward at each end. The behavior consisted of running in one direction (POS, ~5 s), followed by eating at one track end (~10 s), then running in the other direction (NEG, ~5 s), followed by eating at the other track end (~10 s), then back to a POS lap, and so on. Both immediately before (PRE) and after (POST) this behavior (RUN), the rats slept in a small enclosure located away from the track. During all three periods (PRE, RUN, and POST), we continuously recorded each rat's position and head direction, the spiking activity of many individual hippocampal CA1 pyramidal cells (Figure 1), and the local field potentials (local EEG) in the region around these cells. (See Experimental Procedures for details and Table 1 for information on each rat.)

Each rat's spatial experience during RUN was divided into two separate place field sequences (POS, NEG), one for each direction along the track. These sequences (each lasting about 5 s) were determined by ordering the peaks of the cells' smoothed firing rate fields (Figures 2A, 2B, and 2G). In linear environments, a large fraction of place cells fire in only one direction; thus, the two sequences generally consist of different sets of cells. Some cells fired in both directions and thus participated in both sequences. Place fields of such cells were determined separately for each direction. The average number of cells per sequence was 8.5 ± 2.1 . (See Experimental Procedures for details.)

¹Correspondence: albert@cortical.mit.edu (A.K.L.), wilson@ai.mit.edu (M.A.W.)

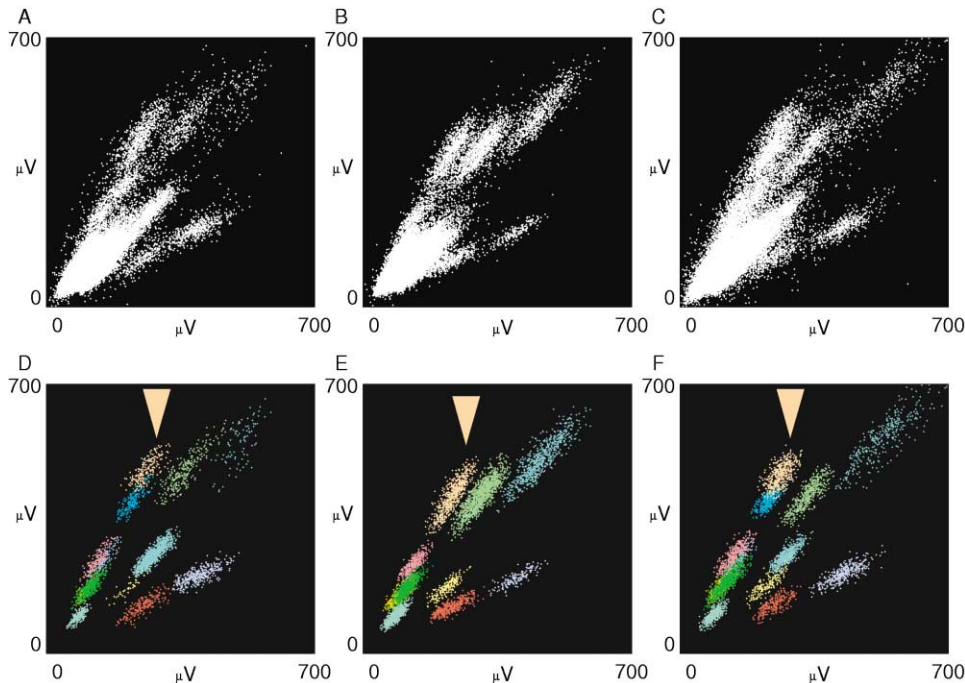


Figure 1. Recording of Spike Data from Individual CA1 Neurons

(A–C) Example of raw data from one tetrode (of RAT1) from which the spikes of several individual CA1 neurons are then clustered. Each point represents the peak amplitude (in microvolts) of a triggered waveform on two of the four tetrode channels (each tetrode consists of four electrodes bundled closely together). All triggered points shown; no noise has been filtered out yet. (D)–(F) shows clusters from (A)–(C), respectively. Clustering is done using waveforms from all four channels. Clarity and stability of clusters across the nearly 4 hr of PRE (A and D, 78 min), RUN (B and E, 23 min), and POST (C and F, 112 min) allow unambiguous identification of the spikes of individual neurons across these three periods. Arrow points to cell 6 in Figures 2A–2F.

Figure 3A illustrates the characteristic activity of a population of CA1 pyramidal neurons during SWS (in this case from POST): short bursts of spikes from several cells with adjacent bursts separated by larger gaps with little firing. Closer inspection of one burst (Figure 2C) reveals that six cells (2,5,7,8,9, and 10) fired in the same order as during RUN (Figure 2B) and were highly compressed in time relative to RUN (2.4 s in RUN behavior becomes 120 ms in SWS). Note that the order matches only if we count the first spike from each cell, with the later spikes representing a continuing burst of activity related (but not limited) to complex spiking of CA1 pyramidal cells (Ranck, 1973). POST SWS contains many such examples of short bursts which closely match the order of their place field peaks in RUN (Figures 2C–2F match Figure 2B, Figures 2H–2J match Figure 2G).

To quantify these sequential patterns, we developed a two-step decoding procedure. First, each SWS population burst of cells from a given RUN sequence was parsed into a single word representing only the relative firing order of the cells within that burst. Then, we computed the probability that the relative firing order within each word would by chance alone match the RUN sequence as well as it actually did. We computed the overall degree of matching between the RUN sequence and (PRE or POST) SWS activity by considering the probabilities from all the words. The details of this procedure are as follows.

To extract the individual population bursts, our parsing procedure consisted of two parameters applied in

order to the activity of a set of cells in SWS. First, for each cell, each set of consecutive spikes with interspike intervals less than max_isi (ms) was grouped into a single event (represented by a single letter identifying that cell) occurring at the time of the first spike. Each isolated spike (i.e. a spike separated by at least max_isi from any other spike of that cell) also produced a letter. Then, the letters from all the cells were merged to form a single, time-ordered string. This string was then broken between every pair of adjacent letters separated by more than max_gap (ms) into a set of shorter strings (words). Each word thus represents the relative order of cell activity within an SWS population burst. We parsed the SWS activity of the cells from the POS and NEG place field sequences separately—the SWS activity of the cells active in the POS behavioral sequence gave one set of words (to be tested for matches to the POS experience), and those from NEG another set (to be tested for matches to the NEG experience).

To quantify the degree of matching between (PRE or POST) SWS activity and a given RUN sequence, we classified the words parsed from SWS into three different groups based on their complexity—pairs, two-letter words with two distinct letters; triplets, three-letter words with three distinct letters; and low-probability trials, defined precisely below, but generally consisting of words with four or more distinct letters as well as any additional repeated letters. Sequence matching within each group was evaluated separately. For pairs, we computed the ratio of the number of pairs in which the

Table 1. Individual Data for Each of the Three Sessions (One for Each of Three Rats)

	RAT1	RAT2	RAT3
Recording duration (min)			
PRE	78	55	37
RUN	23	24	19
POST	112	51	31
Amount of sleep (min)			
PRE	16	45	24
POST	50	36	20
Amount of SWS (min)			
PRE	15	35	22
POST	47	33	19
Track length (cm)	450	180	180
Number of POS laps	30	23	35
Number of NEG laps	29	23	36
Number of cells in POS seq	10	8	5
Number of cells in NEG seq	11	9	8
PRE SWS PA + T + LP trials/min	16.1	32.4	15.6
POST SWS PA + T + LP trials/min	10.9	32.9	9.7
PRE SWS Z scores			
Pair	0.1	-1.5	-1.2
Triplet	-1.9	-1.0	-0.7
Low-probability	-0.9	-0.1	-0.3
POST SWS Z scores			
Pair	0.0	1.3	1.4
Triplet	0.9	1.4	3.4
Low-probability	4.4	4.8	4.7
POST SWS low-probability p value	<3E-4	<6E-5	<4E-3

All of the pair (PA), triplet (T), and low-probability (LP) trial results are for the case of max_isi = 50 ms and max_gap = 100 ms. Maximum interval between PRE, RUN, and POST was 3 min.

two letters matched their order in the RUN sequence (pair matches) to the total number of pairs regardless of letter order (pair trials). The ratio's significance is the number of standard errors above the expected ratio (1/2), assuming either ordering of letters was equally likely, i.e., a Z score. The same was done for triplets, where the expected ratio of three-letter matches (all three letters in order) to all triplets is 1/6, assuming each of the 3! = six possible orderings of three letters was equally likely.

The most important group was the low-probability trials, since it potentially contained long, very low probability matches (four or more letters in order) to the RUN sequence. To analyze this group, we developed a novel combinatorial method that allowed us to quantify the degree of matching between the RUN sequence and the relative order of letters in words of arbitrary length and letter composition. The method involves identifying the best match within a word to the given RUN sequence, then computing the probability that this word would contain such a match or better, assuming all possible permutations of its constituent letters were equally likely. Thus, the probability indicates how likely it is by chance that the order of letters in a word matches the RUN sequence as well as it does. For example, given the RUN sequence 123456789A (where A represents cell 10) (Figure 2B), the best match found in the seven-letter word 325789A (Figure 2C) is a (6,0) match (six letters in order with zero interruptions—here 25789A). Out of the 7! = 5040 possible permutations of the seven letters, there are 12 permutations whose best match is a (6,0) match (not necessarily 25789A), and one permutation with a better match, the perfect (7,0) match 235789A.

Thus, the probability of getting a (6,0) match or better by chance is very small: $(12 + 1) / 7! = 0.0026$ (Figure 3B). This method provides a unified way of computing the probabilities of different matches in words with differing lengths and letter compositions, including matches with interruptions (e.g., the (6,1) match in 246579A, Figure 2E) and matches in words with repeated letters (e.g., the (5,0) match in 22569A8, Figure 2D). (See Experimental Procedures for more details.)

The precise definition of a low-probability trial is a word for which the best possible sequence match produced by optimally rearranging its letters (in general, a (K,0) match for a word with K distinct letters) has probability $\leq P$ of occurring (assuming all permutations are equally likely), i.e., a word which could have contained a match of probability P or rarer based on its constituent letters. A word is a low-probability match if the probability of getting the best match actually found in that word or better is $\leq P$. The expected match/trial ratio is P, with significance determined as before. We chose $P = 1/24 = 1/4!$, the probability of a perfect four-letter match. Thus, pairs cover two-letter matches; triplets, three-letter matches; and low-probability trials, n letter matches for $n \geq 4$. However, P could be any low value (e.g., $P = 0.01$ focuses on n-letter matches for $n \geq 5$). Therefore, our match probability method can analyze sequential firing patterns involving any number of neurons. It is general and can be applied to other types of neural data. It differs from Abeles' method by analyzing the relative ordering of activity as opposed to patterns defined by a set of specific interspike intervals (Abeles and Gerstein, 1988; Abeles, 1991; Abeles et al., 1993).

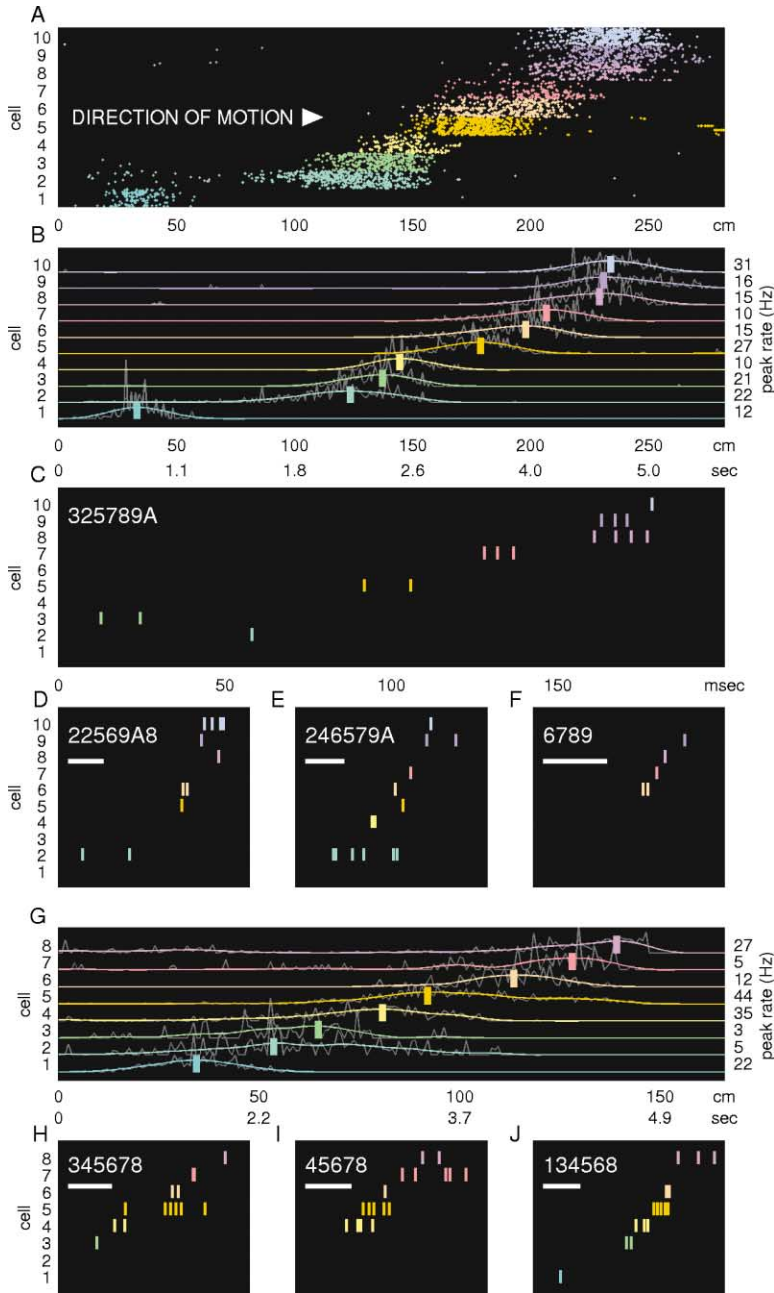


Figure 2. Example Sequences from Behavior (RUN) and Subsequent Sleep (POST)

(A–B) Determination of POS spatial sequence experienced by RAT1 in RUN. (A) Lap-by-lap rasters of all ten cells that had place fields in POS direction laps (i.e., rat running in direction of increasing position values). For each cell, laps 1–30 are stacked from bottom to top. (B) Smoothed place fields (colored lines) of these ten cells. Vertical bars mark the positions of the peaks of the smoothed fields. Smoothed firing rate (Hz) at these peaks shown to the right. Nonuniform time axis below shows time within average lap when above positions were passed. (C) A population burst from RAT1 POST SWS, showing six cells in a row firing in the same order as the POS sequence from RUN (B). Note difference in timescale. (D–F) More examples of RAT1 POST SWS population bursts that match the RUN POS sequence. (G) Same as (B), except for RAT2 POS (rat running in direction of increasing position values). (H–J) RAT2 POST SWS population bursts that match the RUN POS sequence (G). Words extracted from activity in (C)–(F) and (H)–(J) using $\text{max_isi} = 50$ ms and $\text{max_gap} = 100$ ms in upper left corner of each panel (with cell 10 represented in words by the letter A). Bar = 50 ms.

The final results consist of a match/trial ratio and Z score for pairs, triplets, and low-probability words in PRE and POST SWS. This analysis was repeated for a range of max_isi and max_gap values (with the constraint that $\text{max_isi} \leq \text{max_gap}$, since $\text{max_isi} > \text{max_gap}$ would have allowed a single letter to represent activity that extended into the following word) (Figure 4) (note that $Z < 0$ [$Z > 0$]) corresponds to fewer [more] matches than expected based on chance). PRE SWS activity shows no significant similarity to RUN sequences for all parameter values, while POST SWS activity shows significant similarity for a wide range of values. This suggests that the sequential spatial experience was encoded during RUN. Furthermore, the most significant matching in POST occurs for longer, lower-

probability matches (the exception is $\text{max_isi} = 0$. It results in very little POST low-probability matching because if multispikes from individual cells are treated as multiple letters, they tend to interrupt sequence matches). The peak low-probability match/trial ratio for POST ($35/270 = 0.13$, i.e., 35 matches out of 270 words that could have had a low-probability match compared to $270/24 = 11$ expected matches, $Z \geq 7.2$, $p < 4E-9$, Figure 4B) occurs around $\text{max_isi} = 50$ ms and $\text{max_gap} = 100$ ms, suggesting that these may be the best values for decoding CA1 pyramidal cell activity in SWS. These values parse the activity into words such that the resulting mean interword interval is 804 ms, a number consistent with the hippocampal EEG sharp-wave/ripple occurrence rate of approximately 0.5–1 Hz

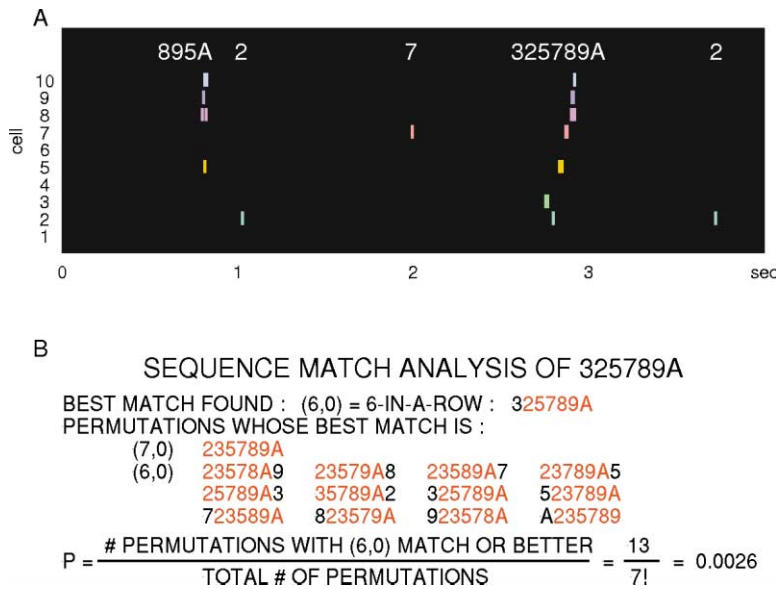


Figure 3. Parsing and Match Probability Method

(A) POST SWS activity of the ten cells from RAT1 RUN POS sequence (Figure 2B). This illustrates the characteristic activity of CA1 pyramidal cells in SWS: short population bursts separated by larger gaps with little activity. Words extracted using $\text{max_isi} = 50$ ms and $\text{max_gap} = 100$ ms shown above each burst. The activity below “325789A” is shown at an expanded time scale in Figure 2C. (B) Example probability calculation for matching of word 325789A (Figure 2C) to RUN sequence 123456789A (Figure 2B). Letters in matches are in red.

during SWS. Significant POST low-probability matching, as well as the other trends of Figure 4, are present in the individual rats (Table 1). All results that follow are pooled over all six RUN sequences of the three rats (two sequences per rat: POS, NEG), using $\text{max_isi} = 50$ ms and $\text{max_gap} = 100$ ms.

Figure 5A shows the match/trial ratios for pairs, triplets, and low-probability trials. The most dramatic deviation from expected occurs for POST low-probability trials: the actual ratio of 0.13 is triple the expected ratio $1/24 (= 0.042)$ ($Z \geq 7.2$, $p < 4E-9$), while the actual ratio for pairs is only 0.52 compared to the expected ratio 0.5 ($Z = 1.6$, $p > 0.05$). Histograms of all words containing matches of probability ≤ 0.1 (Figure 5B) show that POST has many events of probability far less than $1/24$ (and thus the lower bound Z of 7.2 is an even greater underestimate), while PRE shows no evidence of RUN sequence structure. Considering even lower probability events ($P = 0.01$ instead of $1/24 [= 0.042]$) yields an even greater deviation from the expected match/trial ratio in POST (11 matches out of 140 trials, $140 \times 0.01 = 1.4$ expected, $p < 3E-7$).

We performed several control analyses to investigate what might be responsible for the high occurrence of long sequences in POST that match the place field sequences from RUN.

First, we computed the significance for low-probability trials with respect to randomly shuffled versions of the actual RUN sequences. Figure 5C shows that POST specifically matches the determined RUN sequences and not other random sequences. If general non-RUN-sequence effects were responsible for significant matching between POST and the RUN sequences, significant matching should also have occurred for shuffled sequences, but it did not. Thus, the RUN sequence is a special sequence for POST SWS. This shuffle analysis also allows us to test the validity of our match probability method. It shows that our match probability method produces a statistic (i.e., the low-probability Z score) that accurately assesses the significance of sequence

structure within real population activity. In particular, the POST low-probability Z score of 7.2 produced by our method agrees with the Z score of 7.9 determined from the POST shuffle distribution (POST shuffle distribution mean Z score = -0.38 , standard deviation 0.96; thus the POST Z score of 7.2 is $(7.2 - (-0.38)) / 0.96 = 7.9$ standard deviations above the shuffle mean) and likewise for PRE (PRE shuffle distribution mean Z score = -0.31 , standard deviation 0.84; thus the PRE Z score of -0.5 is 0.2 standard deviations below the shuffle mean). Furthermore, the PRE and POST shuffle distributions closely match the distributions predicted by a model in which each low-probability trial is an independent trial with equally likely permutations (Figure 5C). A similar analysis of triplets gives the same result. Thus, the real data are consistent with the equally likely permutation model on which our match probability (and resulting match significance) calculations are based. A nonparametric comparison which uses only match rankings and does not deal with permutations or match probabilities confirms that POST matches RUN sequences more than PRE does ($p < 0.01$, K-S test).

Second, we checked that the long sequence (low probability) matching effect in POST was not simply due to a tendency for cells with nearby place fields to fire together regardless of order (Wilson and McNaughton, 1994). If this were the case, then we would also expect significant matching in POST to the reverse RUN place field sequences (e.g., A987654321 from Figure 2B), but this was not so ($Z \leq 1.1$, $p > 0.16$).

Third, we recomputed the probabilities for POST low-probability trials, this time assuming that all permutations were not equally likely, but rather were weighted by the actual pair ratio in POST (which, at 0.52, was slightly biased in favor of the RUN sequence) (see Experimental Procedures). The idea was to test whether the observed tendency for pairs of letters in POST to occur in the same order as in the RUN sequence could alone explain the unexpectedly large number of low-probability matches. The significance was reduced (as expected)

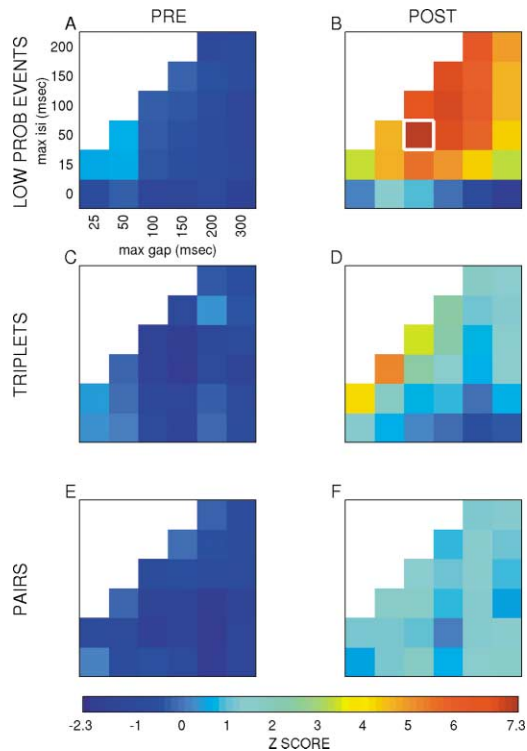


Figure 4. Significance (Z Score) of Matching of PRE and POST SWS Activity to RUN Spatial Sequences as a Function of the Two Parsing Parameters Max_isi and Max_gap, Pooled over All Three Rats

These results show that PRE SWS bursts have no similarity to the RUN sequences for all parsing parameter values, while POST SWS bursts exhibit significant similarity to the RUN sequences. This suggests that the sequential spatial experience was encoded during RUN. POST similarity is most significant in the longer bursts (i.e., low-probability trials), and this is the case for a wide range of parameter values. Absence of sequence similarity in POST for max_isi = 0 shows that it is necessary to treat multispike bursts of a cell as single events in order to decode sequential activity in SWS. White box in (B) indicates maximum POST low-probability match/trial ratio and Z score. This occurs at max_isi = 50 ms and max_gap = 100 ms.

but still highly significant (actual match/trial ratio = 0.11, expected ratio ≤ 0.054 , $Z \geq 4.2$, $p < 2E-4$). Thus, pair biases cannot account for the high occurrence of long, low-probability sequence matches.

Fourth, since the parsing method makes it harder for repeated (versus distinct) letters to appear close together in a given word, we recomputed the match probabilities to account for this. In particular, we did not count permutations in which repeated letters occurred closer together than in the observed word, while the remaining permutations were assumed equally likely. Results were unchanged (low-probability trials: PRE $Z \leq -0.07$, $p > 0.6$, POST $Z \geq 7.5$, $p < 2E-9$).

Fifth, to specifically control for artifactual sequence effects that could result from imperfect clustering (Quirk and Wilson, 1999), we reparsed and recomputed the results after each of the following manipulations. We first eliminated a given fraction of spikes from either the outer border or the low-amplitude tail of each cell's cluster, then either stopped there or further eliminated all remaining spikes that potentially could still be part of a single complex-spike incorrectly split across multiple

cells. For every manipulation, the remaining spikes still contained highly significant low-probability matching in POST (Figure 5D).

Thus, none of these potential causes can account for the high occurrence of long RUN sequence matches in POST.

Furthermore, we computed the significance for low-probability trials with respect to RUN sequences that bridge the 10 s break between POS and NEG direction laps. These wraparound sequences (two per rat) were constructed by adding the first half of the NEG sequence to the end of the last half of the POS sequence (giving one sequence) and by adding the first half of the POS sequence to the end of the last half of the NEG sequence (giving the other sequence). The lack of matching in POST ($Z \leq -1.2$, $p > 0.9$) supports the idea that the POS and NEG sequences represent distinct bursts of experience that are encoded separately. Cells which participate in both POS and NEG sequences do in fact appear in the appropriate positions in both POS- and NEG-like SWS sequences (e.g., RAT1 POS cells 5, 9, 10 in Figures 2C–2F correspond to RAT1 NEG cells 7, 4, 5, respectively, in Figure 6C).

For one rat (RAT2), RUN constituted the first exposure to that environment, while the other two rats had some previous training in their respective RUN environments (though all three rats had been trained before on various linear tracks in other rooms). In both cases (novel, familiar), PRE SWS has no sequence structure related to RUN, while POST does (low-probability trials, novel: PRE $Z \leq -0.1$, $p > 0.6$, POST $Z \geq 4.8$, $p < 6E-5$; familiar: PRE $Z \leq -0.9$, $p > 0.9$, POST $Z \geq 5.5$, $p < 2E-5$; also see Table 1). This suggests that long sequences can be encoded in a single RUN session and that the mechanisms leading to expression of RUN sequences in SWS (possibly including modified synaptic connections) are reset sometime after POST but before the next exposure (which occurred at least a day later; see Experimental Procedures). Similar resetting of other experience-dependent changes in the hippocampus has been observed before (Mehta et al., 1997, 2000).

Figure 6A shows the distribution of durations (time between first and last letters) of all PRE and POST low-probability trials (mean 154 ± 84 ms), as well as all POST low-probability matches (time between first and last letters in match) (mean 106 ± 38 ms). The population bursts that constitute the low-probability trials tend to occur in periods with heightened high-frequency ripple (100–250 Hz, 25–100 ms duration events) activity (Figures 6B and 6C), a prominent feature of hippocampal EEG during SWS. This is not surprising, since it is known that overall CA1 pyramidal cell activity in SWS is greatly increased during such ripple events.

As seen in Figure 2, POST SWS sequence matches are compressed in time relative to the experienced RUN sequence. We estimated the approximate compression factor (CF) by comparing the times of the letters in each POST low-probability match to the times of their place field peaks in the RUN sequence (e.g., Figures 2B and 2G) (see Experimental Procedures). The median of 19.7 (Figure 7A) indicates that sequences in POST SWS occurred approximately 20 times faster than the original experience. To verify this using a different method, we computed the average overlap between the POST low-

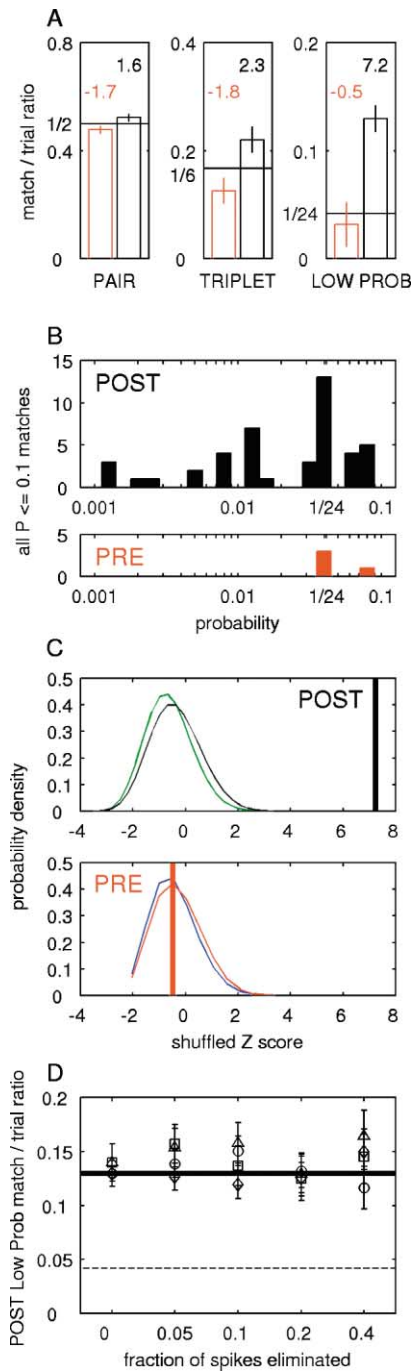


Figure 5. Results Pooled from All Three Rats Using Parameter Values Max_{isi} = 50 ms and Max_{gap} = 100 ms
(A) Actual (bars) versus expected (horizontal lines) match/trial ratio for pairs, triplets, and low-probability trials for PRE (left, red) and POST (right, black) SWS. Vertical lines: ± 1 standard error. Z scores above. Actual match/trial numbers (PRE, POST): pairs (654/1371, 655/1255), triplets (32/255, 57/259), low-probability trials (3/95, 35/270). (B) All PRE and POST SWS matches to RUN sequence of probability ≤ 0.1 . Log scale on abscissa. (C) Distributions of PRE (below, red) and POST (above, black) SWS low-probability Z scores with respect to randomly shuffled RUN sequences. Thus, high Z scores do not occur by chance. Vertical lines: PRE (red) and POST (black) Z scores for actual RUN sequences. Predicted distributions for PRE (below, blue) and POST (above, green) assuming each low-probability trial is an independent trial with equally likely permuta-




Figure 6. Low-Probability Trials and Ripple Events
(A) Duration of all low-probability trials (line) in PRE and POST SWS, and matches (solid) in POST SWS. (B) Timing of ripples (midpoint) with respect to PRE and POST SWS low-probability trials (center of mass time of letters) (for RAT1 and RAT2 only). Binsize = 20 ms. (C) Example showing timing of ripples about a POST SWS low-probability match (from RAT1 NEG).

probability matches and various time-compressed templates of the overall place field activity in RUN (e.g., the smoothed firing rates in Figures 2B and 2G as a function of time) (see Experimental Procedures). The template overlap method detects broadly matching temporal activity patterns at any selected time scale (CF) despite numerous extra spikes and imperfect ordering. The peak occurs around CF = 16 (Figure 7B), in agreement with the first method.

Finally, we looked for evidence of any additional RUN sequence structure at any time scale (CF) in POST SWS besides that detected using the parsing and match

tions. (D) Control for robustness of POST SWS low-probability sequence matching effect with respect to imperfect clustering. Match/trial ratios ± 1 standard error after systematic elimination of particular spikes (see text). Solid and dotted horizontal lines represent original and expected (1/24) POST low-probability match/trial ratios, respectively. Sequence effect highly significant under every manipulation.

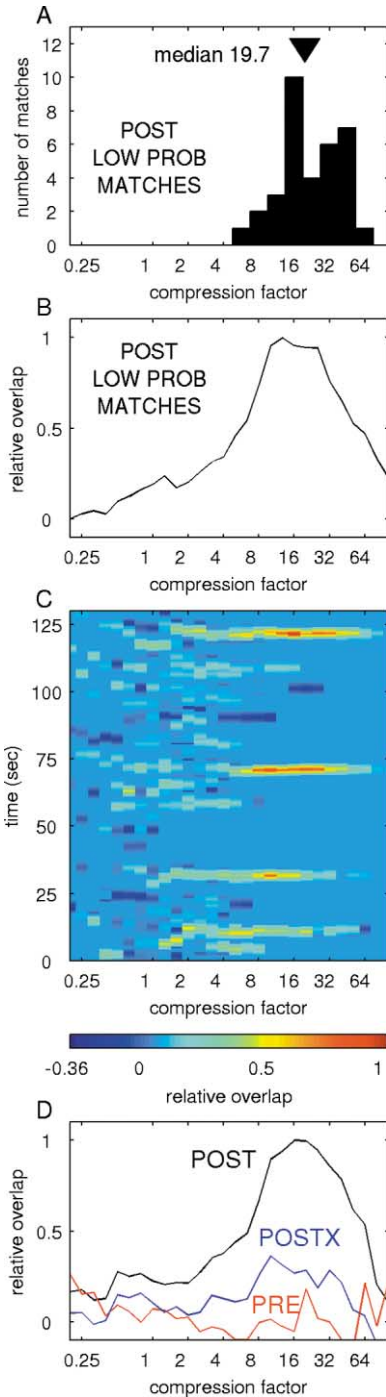


Figure 7. Time Compression of Sequences

(A and B) CF of POST SWS low-probability match sequences with respect to RUN behavioral sequences (e.g., Figures 2B and 2G in terms of time). CF > 1 means SWS sequences are compressed in time relative to behavioral sequences. (A) Distribution of CFs for individual POST low-probability matches; median CF = 19.7. (B) Average overlap between behavioral sequences and POST low-probability matches as a function of CF; peak occurs around CF = 16. (C) Example showing overlap between POS sequence behavioral template (Figure 2G in terms of time) and segment of POST SWS for RAT2 as a function of time. Unlike in (B), the overlap here is computed continuously in time, not just around low-probability matches. Red horizontal lines indicate significant matching which

probability method. We again used the template method because of its ability to detect broadly matching activity at any time scale. For each POS and NEG RUN template, we computed the overlap as a function of time for many different CFs (Figure 7C). Figure 7D shows the overlap averaged over all of PRE SWS, all of POST SWS, and all of POST SWS except for those times immediately around sequence matches detected using our parsing and match probability method (see Experimental Procedures). We find no evidence of any sequence structure in PRE SWS. POST SWS has sequence structure with a single peak around CF = 20, and this peak disappears when the small fraction of overlap scores around matches of probability $\leq 1/6$ (i.e., triplet matches or longer) is eliminated. Thus, the sequences found with our decoding method (using $\text{max_isi} = 50$ ms and $\text{max_gap} = 100$ ms) likely represent all the RUN sequence structure present in POST SWS (though there could be additional non-RUN-sequence structure in PRE and/or POST SWS).

Discussion

In general, not all differences between PRE and POST can be attributed to learning. RUN may result in many changes (e.g., increases in firing rate) that are not related to the particular spatial sequences experienced. However, we demonstrate that random sequences do not result in POST matching (Figure 5C). Thus, the observed matching between POST and the actual RUN sequences cannot simply be explained by phenomena such as general changes in firing rates (which would have affected random sequences similarly). Furthermore, we demonstrate that reverse RUN sequences do not result in POST matching. Thus, the matching between POST and the RUN sequences cannot be explained by non-sequence-specific increases in the coactivity of cells with nearby place fields. Therefore, RUN sequence structure is specifically present in POST. Moreover, we demonstrate that the longer sequences in POST cannot be accounted for by the occasional, random chaining together of shorter pair sequences. In contrast, PRE shows no evidence of RUN sequence structure of any length (from pairs to longer sequences) or at any time scale beyond that expected by chance (whether chance is based on the assumption of equally likely permutations or on the distribution of matching of random sequences to PRE) (Figures 4, 5A–5C, and 7D). Therefore, we conclude that long spatial sequences experienced during behavior (i.e., the separate POS and NEG place field sequences) are indeed encoded precisely (i.e., preserve relative firing order) and rapidly (i.e., within a single RUN session).

occurs intermittently. Note that (D) is essentially created by collapsing these results (plus all of the other SWS data) vertically. (D) Overlap as function of CF averaged over all of PRE SWS, all of POST SWS, and POSTX (all of POST SWS except for those times immediately around matches of probability $\leq 1/6$). This shows that the only significant RUN sequence structure found in PRE or POST SWS occurs in POST at CFs of around 20 and that this structure can be fully accounted for by the sequences found using our parsing and match probability method. All matches referred to in this figure extracted using $\text{max_isi} = 50$ ms and $\text{max_gap} = 100$ ms.

These memories can be decoded by applying our two-parameter parsing and match probability method to the population activity of CA1 pyramidal neurons during SWS. Proper decoding requires that the activity of multispike bursts from individual neurons be treated as single events. (Thus, our results suggest that these first spikes, determined with `max_isis` as large as 50 ms, might be particularly important for indicating information about temporal order.) The intermittent, time-compressed sequences found with this method likely represent all the RUN sequence structure present in POST SWS (Figure 7D).

Our combinatorial match probability method for analyzing relative firing order among an arbitrary number of neurons is general and can be applied to other types of neural data (and it is not just limited to spike data, e.g., it can be applied to EEG events from multiple locations). All that is necessary is a preliminary parsing step appropriate for each type of data. Here, since hippocampal CA1 pyramidal cell activity in SWS consists of brief, intermittent population bursts, each burst naturally provided a word. In the case of an experiment which consists of a series of trials, each trial might provide a word.

Previous work has shown traces of hippocampal activity from RUN in sleep using various methods (Pavlidis and Winson, 1989; Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996; Kudrimoti et al., 1999; Nádasdy et al., 1999; Louie and Wilson, 2001). Skaggs and McNaughton, whose analysis was limited to pair interactions, showed a tendency for the average firing order between pairs of cells to be conserved from RUN to POST and also (but less so) from PRE to RUN, although they did not discuss possible timescale differences between RUN and sleep (Skaggs and McNaughton, 1996). Here, we find that the POST pairwise sequence effect is small (0.52 versus an expected 0.5, which is similar in magnitude to that found using their method) and that it cannot explain the much stronger effect (>3 times expected) of the high occurrence of longer sequences. Using a nonspatial running wheel task, Nádasdy et al. (1999) showed that many triplet spike sequences were common to PRE, RUN, and POST. Unlike our place field sequences, their RUN spike sequences were not clearly related to any specific behavioral events. Furthermore, while they noted that their sleep sequences generally occurred at a faster timescale than the RUN sequences, no compression factor was computed. Our work differs from both of these studies by analyzing longer sequences (i.e., lengths of four and more versus pairs and triplets). In addition, we demonstrate that the specific sequences experienced in RUN are strongly present in POST but completely absent (i.e., match/trial ratios below expected based on chance, i.e., Z scores < 0 , Figure 5A) in PRE, thus implying the RUN sequences were rapidly encoded during that RUN itself. In contrast, the other studies reported either a smaller POST effect along with a small amount of RUN structure in PRE (Skaggs and McNaughton, 1996) or a larger POST effect along with a larger amount of RUN structure in PRE (Nádasdy et al., 1999). Strong, specific matching in POST combined with no matching in PRE makes it unlikely that we are merely observing intrinsic features of the network conserved across PRE-RUN-POST and unrelated to learning in RUN.

What could be the mechanisms by which these RUN sequences are encoded? Known plasticity rules require precisely ordered activity occurring within short time windows (Levy and Steward, 1983; Gustafsson et al., 1987; Markram et al., 1997; Bi and Poo, 1998). Presynaptic activity must precede postsynaptic activity by less than approximately 20 ms to yield strengthening of connections, while the reverse order yields weakening. To encode spatial sequences experienced at the behavioral timescale (1 s) using plasticity mechanisms that operate at this shorter timescale (10 ms), hippocampal phase precession has been proposed as a possible bridging mechanism (O'Keefe and Recce, 1993; Blum and Abbott, 1996; Skaggs et al., 1996; Mehta et al., 2002). Connections may be modified within the hippocampal CA3 network (Blum and Abbott, 1996; Wallenstein and Hasselmo, 1997), with encoded sequences being replayed by triggering a cascade of activity (August and Levy, 1999). Alternatively, the same changes in the CA3-CA1 feedforward network that produce asymmetric place fields (Mehta et al., 2000) may be involved. Asymmetric drive translated into latency could both encode (Mehta et al., 2002) and replay RUN sequences. There is evidence that the sharp-wave/ripple events, which generally accompany SWS population bursts (Figures 6B and 6C), are generated by the hippocampus itself (Buzsáki et al., 1983). If so, the SWS sequences we observe might predominantly reflect the intrinsic connectivity of the hippocampus, thus suggesting the sequential experience was indeed encoded within the hippocampus (versus in external structures only).

Recently, long stretches (up to 2 min) of RUN sequential structure have been observed during rapid eye movement (REM) sleep at timescales similar to the original experience (CF = 0.7) (Louie and Wilson, 2001). While the highly compressed, 100 ms timescale SWS sequences could be replayed via intrinsic hippocampal mechanisms, the extended replay in REM likely involves activation of a broader network that can sustain regular activity for several minutes. Unlike the SWS sequences, REM replay was prominent in PRE, representing spatial experience encoded at least a day before. The SWS sequences we find immediately after only one RUN session may represent rapid learning of discrete episodes (POS, NEG) by the hippocampus, while REM replay could represent a later stage of memory processing involving longer stretches of experience and additional brain areas.

Sleep may play an important role in learning (Buzsáki, 1989; Pavlidis and Winson, 1989; Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996; Kudrimoti et al., 1999; Nádasdy et al., 1999; Dave and Margoliash, 2000; Stickgold et al., 2000; Frank et al., 2001; Louie and Wilson, 2001). In addition to revealing asymmetric modifications to the hippocampal synaptic weight matrix, replay of sequences learned in behavior may indicate a process of memory consolidation. The occurrence of low-probability matches cannot be accounted for by pairwise biases in firing order, indicating that hippocampal replay consists of brief bursts of long sequences that occur intermittently rather than a more continuous replay of pairwise biases. These intermittent, time-compressed sequences may represent discrete packets of learned information being broadcast by the

hippocampus to modify target structures elsewhere in the brain (Buzsáki, 1989; Pavlides and Winson, 1989; Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996; Siapas and Wilson, 1998; Kudrimoti et al., 1999; Nádasdy et al., 1999), with the precise ordering and timescale (10 ms) required by plasticity mechanisms (Levy and Steward, 1983; Gustafsson et al., 1987; Markram et al., 1997; Bi and Poo, 1998).

Experimental Procedures

Electrophysiology

Spikes from many individual CA1 pyramidal neurons (RAT1, 57 cells; RAT2, 49 cells; RAT3, 26 cells) were recorded simultaneously from each rat (adult male Long Evans) using a microdrive array containing four to eight independently adjustable tetrodes and placed above the right dorsal hippocampus (3.6 mm posterior, 2.4 mm lateral with respect to bregma). Tetrodes were first lowered over 1–2 weeks to the CA1 pyramidal cell layer. Then fine adjustments to maximize the number of isolatable pyramidal cells on each tetrode were made with net movement limited to 20–120 μm per day. For stability, recordings were made no less than 12 hr after the last adjustment. Individual cells were isolated from each tetrode (Figure 1) via spike waveform clustering using XCLUST (M.A.W.) (Wilson and McNaughton, 1993). Spike times were recorded with precision 0.1 ms. Rat head position (resolution 0.8 cm) and head direction were sampled at 30 Hz. Hippocampal CA1 pyramidal cell layer EEG (filtered between 1–475 Hz and sampled at 1.5 kHz) was recorded from each tetrode. All procedures were performed in accordance with institutional guidelines.

Behavioral Details

For several days before recording, rats were trained on various linear tracks to run back and forth (without stopping in the middle of the track) for chocolate reward at each end. The last training session occurred no less than 1 day before recording. For the novel rat (RAT2), RUN constituted the first exposure to that environment. Of the two familiar rats, for RAT1 the immediately preceding exposure to the RUN environment occurred 5 days before (with no intervening training in any environment), while for RAT3 it occurred 1 day before (along with an intervening training session in a different environment on that same preceding day). For PRE and POST, rats were left in a 50(l) \times 50(w) \times 75(h) (cm) black box, open at the top and lighted from above. Rats were placed on a 20 cm diameter padded circular dish centered and raised 40 cm above the box floor so that movement was restricted and the side walls could not be reached. There the rats' behavior consisted of sleeping, grooming, or remaining still while awake. After PRE, the rat was immediately moved to the track for RUN, then immediately returned to the box for POST.

Determination of RUN Place Field Sequences

Only a fraction of CA1 pyramidal cells have place fields in any given environment (e.g., a track). To select the cells used to determine the POS and NEG direction place field sequences, we applied a set of criteria to each cell twice, once for each direction. Thus, each cell could be in 0, 1, or both sequences. Specifically, to determine which cells would be included in a given direction's sequence, we first eliminated all spikes fired at the track ends (because of variable behavior and behavioral state there) or when the rat stopped moving in the middle of the track (which was very infrequent). With the remaining spikes, we took all pyramidal cells firing ≥ 1 spike/lap (mean) in that direction, then eliminated those cells that either exhibited double-peaked place fields within that direction (thus no unique position in sequence) or which stopped firing in that direction for the middle or last 1/3 of the RUN session (indicating unreliable firing). We then computed the smoothed one-dimensional directional firing rate field of the remaining spikes of each remaining cell. The order of peaks was the same whether smoothing with a Gaussian of $\sigma = 5, 10, \text{ or } 15$ cm. For the six directions in the three rats, a total of 64 cells (46 of them distinct) fired ≥ 1 spike/lap, of which four had double-peaked fields and nine stopped firing for the middle and/or last 1/3 of the session, leaving 51 cells (40 distinct). Table 1 shows

the number of these cells in the POS and NEG sequences of each rat. The POS and NEG sequences of a rat were generally not reversed versions of each other because many cells had place fields in only one direction and because those cells with place fields in both directions did not necessarily have both fields in the same location.

Identification of SWS

Sleep, identified based on visual assessment (through videotape) of rat posture during PRE and POST, was partitioned into SWS and REM as follows. REM was identified as periods with an elevated ratio (averaged every 1 s) of hippocampal EEG power in the theta band (5–12 Hz) to overall power (1–475 Hz). The remainder was classified as SWS. All SWS from PRE and POST was used in our analyses.

Match Probability Analysis

A word, W (which may have repeated letters), contains a (x,y) match to a given sequence, S (which cannot have repeated letters), if there exists $x + y$ consecutive letters in W of which at least x letters are in strictly increasing order in S (but not necessarily consecutive in S). For example, if $S = 123456789$, the word 11377 does not contain a $(5,0)$ match but does contain a $(3,0)$ match, and the word 13436892 contains a $(6,1)$ match. We construct an ordered list, L , of matches to S (best to worst) as follows. The list starts with $(K,0)$ as the best match, where $K = \text{the length of } S$. Next are (x,y) such that $x - y = K - 1$, ordered by decreasing x . Thus $(K,1)$ is the second best and $(K - 1,0)$ the third best match. Next are (x,y) such that $x - y = K - 2$, ordered again by decreasing x . This continues for $x - y \geq 2$. Thus $(2,0)$ is the last match in the list. Matches with $x - y < 2$ are not considered. This ordering obeys the obvious constraints that $(x,y) > (x - 1,y) > (x - 2,y) > \dots$ and $(x,y) > (x,y + 1) > (x,y + 2) > \dots$ (where " $>$ " means "better than") and, more generally, obeys the constraint that a word which contains a given match does not necessarily contain a match better than it. This ordering also balances rewarding larger x (longer matches) with penalties for larger y (more interruptions). The best match to S found in a given word W is the best match in L that is found in W . For a given word W of length n letters (not necessarily all distinct), the probability of getting a given match (x,y) from L or better is the fraction of the $n!$ permutations of the letters of W that contain a match (x,y) or better according to L . This provides a unified way of computing the probabilities of different matches in words with differing lengths and letter compositions, including matches with interruptions and matches in words with repeated letters. (A different L in which for any two matches $(x_1,y_1) > (x_2,y_2)$ if (a) $x_1 > x_2$ or if (b) $x_1 = x_2$ and $y_1 < y_2$ gives similar results, e.g., for $\text{max_isi} = 50$ ms and $\text{max_gap} = 100$ ms: low-probability PRE $Z \leq -0.5$, $p > 0.4$; POST $Z \geq 7.8$, $p < 3E-10$. This suggests that our results do not depend on the minute details of how matches are ordered.)

The reason for computing the probability that a word would contain a match as good as or better than the best match found in that word, assuming all possible permutations of its constituent letters were equally likely, is this: it is exactly analogous to the statistical procedure of determining a p value, such as in computing the probability that one would observe 80 heads in 100 tosses of a coin, assuming the coin is fair. To evaluate how unlikely this is by chance, one computes the probability of getting 80 or more heads assuming the null hypothesis of a fair coin. Here, 80 heads corresponds to the best match found in a word, while more than 80 heads corresponds to the better matches.

Determination of Significance

Given M pair matches out of N pair trials, the normal approximation of the binomial distribution gives Z score = $(M - NP) / \sqrt{NP(1 - P)}$ where $P = 1/2$. The same formula applies to triplets (with $P = 1/6$), and low-probability trials (with $P = 1/24$). For low-probability trials, we computed a lower bound Z for POST and upper bound Z for PRE, i.e., we were most conservative in testing for matching in POST and nonmatching in PRE. To calculate p values, we used the normal approximation for pairs (since the deviations from expected values were small) and the exact binomial distribution for low-probability trials (since the deviations were often very large).

Control for Pair Bias

To correct for pair bias, we recomputed match probabilities as follows. Instead of weighting each permutation equally, we weighted each n letter permutation of a word by treating each of its $n(n-1)/2$ letter pairs as an independent trial with bias $B = 0.52 =$ the observed POST pair match/trial ratio. That is, we weighted each permutation of a word by $B^f(1-B)^{H-f}/H$, where $f =$ the total number of letter pairs in the permutation that were in the same order as in RUN, $r =$ the total number of letter pairs in the permutation that were in the opposite order as in RUN, and $H =$ normalization to make the weights of all the permutations sum to one. This gives an upper bound of the effect of the pair bias since not all these pairs are independent. As an example, the probability of a quadruplet match (four-letter word with four distinct letters, all in order) changed from $1/24 = (0.042)$ to 0.054 with this pair bias.

Identification of Ripples

Hippocampal EEG was filtered in the ripple band (100–400 Hz), the mean and standard deviation of the absolute value of this filtered trace for all of SWS was computed, and a threshold was set to five standard deviations above the mean. Threshold crossings (times when the absolute value of the filtered trace was above threshold) with interevent interval ≤ 20 ms were grouped together to form a single interval. Intervals < 20 ms in duration were rejected, and the rest were classified as ripples. Threshold was verified by visual inspection of correspondence between these classified ripples and the original (1–475 Hz) EEG trace. Results in Figure 6B were unchanged for several different values of threshold.

CF Analysis of Figure 7A

Given a low-probability match (x,y) , we took the x letters in the match, then determined the pairwise CF for each of the $x(x-1)/2$ pairs of those letters by dividing the time interval between the pair of letters in that match into the time interval between the corresponding pair of place field peaks in the RUN sequence (e.g., Figures 2B and 2G in terms of time). The CF for the match was taken to be the median of those $x(x-1)/2$ pairwise CF values.

Template Analysis of CFs

Computing template overlap is described with the following example. Figure 2B as a function of time gives a 5.5 s template of RAT1 RUN POS activity. The smoothed firing rate of each of the ten cells was binned into 55 100 ms bins. To compute the overlap with SWS activity at a certain CF, a window of duration (5.5 s) / CF from SWS was binned into ten cells \times 55 (100 ms) / CF bins. For each cell, the number of spikes in each of the 55 bins was smoothed with a Gaussian of $\sigma = 1$ bin. Then the correlation coefficient between the RUN template and SWS window for each of the ten cells was computed, summed, then divided by the number of cells (here, ten). (This process is similar to Louie and Wilson, 2001.)

To compute the overlap around each low-probability match, the overlap was computed with all windows within \pm one window duration (which depended on the CF) of the match's center of mass time, at a resolution of 1/10 of a window duration (i.e., 11 windows total). This accommodated inexact alignment. The overlap score for a given match was the Z score of the maximum of the 11 overlaps with respect to the distribution of the maxima of the 11 overlaps with cell-shuffled versions of the template (i.e., the original template except with the identities of the cells randomly exchanged). The average over all low-probability matches is shown in Figure 7B.

In computing the average overlap over a long segment of time (Figure 7D), we took care to do it in a manner that treated each CF equivalently so that the overlap scores could be compared across different CFs. First, to compute the overlap as a function of time during PRE (POST) SWS for a given CF (Figure 7C), a sliding window (duration = [RUN template duration] / CF) was moved across PRE (POST) SWS activity in increments of 1/10 of a window duration. At every tenth window, the last ten windows were grouped together. If at least one of the ten windows in a group contained a minimum level of activity (here set to four cells), that group was considered valid. The Z score (with respect to cell-shuffled templates) of the maximum overlap for a valid group of windows was the overlap score at that time. The overlap averaged over time for all of PRE

(POST) SWS for that CF was the mean Z score of all valid groups in PRE (POST) SWS. The overlap for all of POST SWS (for a given CF) except for the times around a set of matches M was the mean Z score of all valid groups in POST SWS except for a small fraction of groups: if the window with maximum overlap for a valid group covered the center of mass time of letters of a match from M , then that group's Z score was excluded.

Acknowledgments

This work was supported by RIKEN and NIH (MH-58880) grants to M.A.W. and a Howard Hughes Medical Institute predoctoral fellowship to A.K.L. We are grateful to W. Asaad, W. Bradley, E. Jonas, K. Lee, G. Liu, K. Louie, M. Mehta, E. Miller, S. Seung, S. Vempala, and A. Wagner for valuable discussions.

Received: May 17, 2002

Revised: October 15, 2002

References

- Abeles, M. (1991). *Corticonics: Neural Circuits of the Cerebral Cortex* (New York: Cambridge University Press).
- Abeles, M., and Gerstein, G.L. (1988). Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *J. Neurophysiol.* **60**, 909–924.
- Abeles, M., Bergman, H., Margalit, E., and Vaadia, E. (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. Neurophysiol.* **70**, 1629–1638.
- Aertsen, A.M.H.J., Gerstein, G.L., Habib, M.K., and Palm, G. (1989). Dynamics of neuronal firing correlation: modulation of "effective connectivity". *J. Neurophysiol.* **61**, 900–917.
- August, D.A., and Levy, W.B. (1999). Temporal sequence compression by an integrate-and-fire model of hippocampal area CA3. *J. Comput. Neurosci.* **6**, 71–90.
- Bi, G.Q., and Poo, M.M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* **18**, 10464–10472.
- Blum, K.I., and Abbott, L.F. (1996). A model of spatial map formation in the hippocampus of the rat. *Neural Comput.* **8**, 85–93.
- Buzsáki, G. (1989). Two-stage model of memory trace formation: a role for "noisy" brain states. *Neuroscience* **37**, 551–570.
- Buzsáki, G., Leung, L.W., and Vanderwolf, C.H. (1983). Cellular bases of hippocampal EEG in the behaving rat. *Brain Res.* **287**, 139–171.
- Dave, A.S., and Margoliash, D. (2000). Song replay during sleep and computational rules for sensorimotor vocal learning. *Science* **290**, 812–816.
- Fortin, N.J., Agster, K.L., and Eichenbaum, H.B. (2002). Critical role of the hippocampus in memory for sequences of events. *Nat. Neurosci.* **5**, 458–462.
- Frank, M.G., Issa, N.P., and Stryker, M.P. (2001). Sleep enhances plasticity in the developing visual cortex. *Neuron* **30**, 275–287.
- Gustafsson, B., Wigström, H., Abraham, W.C., and Huang, Y.Y. (1987). Long-term potentiation in the hippocampus using depolarizing current pulses as the conditioning stimulus to single volley synaptic potentials. *J. Neurosci.* **7**, 774–780.
- Kesner, R.P., Gilbert, P.E., and Barua, L.A. (2002). The role of the hippocampus in memory for the temporal order of a sequence of odors. *Behav. Neurosci.* **116**, 286–290.
- 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.
- Levy, W.B., and Steward, O. (1983). Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience* **8**, 791–797.
- Louie, K., and Wilson, M.A. (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* **29**, 145–156.
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997).

- Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–215.
- Mehta, M.R., Barnes, C.A., and McNaughton, B.L. (1997). Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc. Natl. Acad. Sci. USA* 94, 8918–8921.
- Mehta, M.R., Quirk, M.C., and Wilson, M.A. (2000). Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* 25, 707–715.
- Mehta, M.R., Lee, A.K., and Wilson, M.A. (2002). Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* 417, 741–746.
- Meister, M., Wong, R.O., Baylor, D.A., and Shatz, C.J. (1991). Synchronous bursts of action potentials in ganglion cells of the developing mammalian retina. *Science* 252, 939–943.
- Milner, B. (1966). Amnesia following operation on the temporal lobes. In *Amnesia*, C.W.M. Whitty and O.L. Zangwill, eds. (London: Butterworths), pp. 109–133.
- Mishkin, M., Vargha-Khadem, F., and Gadian, D.G. (1998). Commentary: amnesia and the organization of the hippocampal system. *Hippocampus* 8, 212–216.
- Morris, R.G., Garrud, P., Rawlins, J.N., and O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature* 297, 681–683.
- Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., and Buzsáki, 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.
- O'Keefe, J., and Nadel, L. (1978). *The Hippocampus as a Cognitive Map* (Oxford: Clarendon Press).
- O'Keefe, J., and Recce, M. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3, 317–330.
- Pavrides, 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.
- Quirk, M.C., and Wilson, M.A. (1999). Interaction between spike waveform classification and temporal sequence detection. *J. Neurosci. Methods* 94, 41–52.
- Ranck, J.B., Jr. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. I. Behavioral correlates and firing repertoires. *Exp. Neurol.* 41, 461–531.
- Scoville, W.B., and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry* 20, 11–21.
- 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.
- Skaggs, W.E., McNaughton, B.L., Wilson, M.A., and Barnes, C.A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6, 149–172.
- 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.
- Wallenstein, G.V., and Hasselmo, M.E. (1997). GABAergic modulation of hippocampal population activity: sequence learning, place field development, and the phase precession effect. *J. Neurophysiol.* 78, 393–408.
- Wessberg, J., Stambaugh, C.R., Kralik, J.D., Beck, P.D., Laubach, M., Chapin, J.K., Kim, J., Biggs, S.J., Srinivasan, M.A., and Nicolelis, M.A. (2000). Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* 408, 361–365.
- 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.
- Zola-Morgan, S., Squire, L.R., and Amaral, D.G. (1986). Human amnesia and the medial temporal region: enduring memory impairment following a bilateral lesion limited to field CA1 of the hippocampus. *J. Neurosci.* 6, 2950–2967.