

Hippocampal Replay Is Not a Simple Function of Experience

Anoopum S. Gupta,^{1,2,*} Matthijs A.A. van der Meer,³ David S. Touretzky,^{2,4} and A. David Redish^{3,*}

¹Robotics Institute

²Center for the Neural Basis of Cognition

Carnegie Mellon University, Pittsburgh, PA 15213, USA

³Department of Neuroscience, University of Minnesota, Minneapolis, MN 55455, USA

⁴Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA

*Correspondence: anoopum@cmu.edu (A.S.G.), redish@umn.edu (A.D.R.)

DOI 10.1016/j.neuron.2010.01.034

SUMMARY

Replay of behavioral sequences in the hippocampus during sharp wave ripple complexes (SWRs) provides a potential mechanism for memory consolidation and the learning of knowledge structures. Current hypotheses imply that replay should straightforwardly reflect recent experience. However, we find these hypotheses to be incompatible with the content of replay on a task with two distinct behavioral sequences (A and B). We observed forward and backward replay of B even when rats had been performing A for >10 min. Furthermore, replay of nonlocal sequence B occurred more often when B was infrequently experienced. Neither forward nor backward sequences preferentially represented highly experienced trajectories within a session. Additionally, we observed the construction of never-experienced novel-path sequences. These observations challenge the idea that sequence activation during SWRs is a simple replay of recent experience. Instead, replay reflected all physically available trajectories within the environment, suggesting a potential role in active learning and maintenance of the cognitive map.

INTRODUCTION

Individual place cells in the hippocampus fire in particular locations (place fields) within an environment (O'Keefe and Dostrovsky, 1971; O'Keefe and Nadel, 1978; Redish, 1999). The sequential activation of these cells can be viewed as a neural representation of the animal's trajectory. These sequences are repeated or "replayed" during sharp wave ripple complexes (SWRs), a network event observed in the hippocampal local field potential (O'Keefe and Nadel, 1978; Buzsáki et al., 1983), when the animal sleeps after active behavior (Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996; Kudrimoti et al., 1999; Nadasdy et al., 1999; Lee and Wilson, 2002; Sutherland and McNaughton, 2000) and during pausing behavior in the awake state (Foster and Wilson, 2006; Jackson et al., 2006; O'Neill

et al., 2006; Csicsvari et al., 2007; Diba and Buzsáki, 2007; Johnson et al., 2008; Karlsson and Frank, 2009; Davidson et al., 2009).

Replay of behavioral sequences in the hippocampus provides a possible mechanism for several critical functions in which the hippocampal formation plays a role: consolidation of experiences into long-term memory (Marr, 1971; Buzsáki, 1989; Sutherland and McNaughton, 2000), incorporation or consolidation of information into cognitive schemas or cognitive maps (O'Keefe and Nadel, 1978; Tse et al., 2007), and learning and planning future experiences (Foster and Wilson, 2006; Diba and Buzsáki, 2007; Karlsson and Frank, 2009). To understand the role of SWRs and replay in hippocampal function, recent studies have blocked CA3 output (the site of SWR generation; Ylinen et al., 1995) (Nakashiba et al., 2009) and have directly disrupted SWRs after learning, during sleep (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010; J.C. Jackson et al., 2009, Soc. Neurosci., abstract), showing that these manipulations impair performance on hippocampal-dependent behavioral tasks. Although these studies show that SWRs are clearly important for correct task performance, they cannot specify the nature of learning taking place during these events. The information or content expressed by replay determines what is learned (Marr, 1971; Buzsáki, 1989; Wilson and McNaughton, 1994; Redish and Touretzky, 1998; Redish, 1999; Sutherland and McNaughton, 2000; Foster and Wilson, 2006); thus, by studying the content of replay, we can begin to understand the properties of this learning process.

SWRs and associated sequence replay occur both in slow-wave sleep and during awake rest states. During the awake rest state, representations of previous experiences are reactivated (Jensen and Lisman, 2000; Jackson et al., 2006; O'Neill et al., 2006), and are replayed both in the order that they were experienced (forward replay; Diba and Buzsáki, 2007; Johnson et al., 2008; Karlsson and Frank, 2009; Davidson et al., 2009) and in the reverse order that they were experienced (backward replay; Foster and Wilson, 2006; Csicsvari et al., 2007; Diba and Buzsáki, 2007; Davidson et al., 2009).

Models of forward replay suggest that they reflect previous experience and extend forward due to Hebbian learning from those experiences (Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996; Redish et al., 1998; Redish, 1999; Sutherland and McNaughton, 2000). This is supported by evidence that reactivation increased with experience on a track (Jackson et al., 2006; O'Neill et al., 2006) and reflected both the current

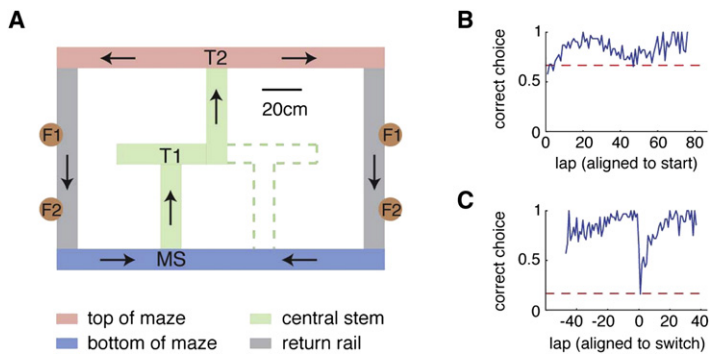


Figure 1. The Two-Choice T Maze

(A) The maze had two possible physical configurations; the second is indicated by dotted green lines. Noteworthy locations on the maze are labeled as follows: maze start (MS), turn 1 (T1), turn 2 (T2), feeder 1 (F1), and feeder 2 (F2). In addition, the task entailed three reward contingencies reflecting a decision made at the second choice point (T2): animals were trained to turn left at the final choice, right at the final choice, or to alternate on a lap-by-lap basis. During recording days, the contingency was changed approximately midway through the task. Place fields did not change between contingencies (Figure S1). (B) Correctness of final choice (T2) aligned to the start of the session. If the task was alternation, the first lap was always deemed correct, so chance performance over all tasks is $2/3 = 66\%$. (C) Correctness of final choice (T2) aligned to the time of contingency switch. Chance of a preswitch correct behavior being a correct behavior after the switch is $1/6 = 16\%$.

and previous environments experienced within a day (Jackson et al., 2006; Karlsson and Frank, 2009). However, others have found preferential reactivation (Cheng and Frank, 2008) and replay (Foster and Wilson, 2006) in novel environments, suggesting that the frequency of replay may be inversely related to experience and therefore the relationship between replay content and experience may not be straightforward. As we will show, detailed identification of the trajectories represented during replay reveals trajectories never actually experienced and depend on the task contingency more than the specific experiences within a day.

Likewise, backward replay was thought to begin at the animal's location in the environment and step backward to replay the immediate past (Foster and Wilson, 2006; Csicsvari et al., 2007), supporting the idea that backward replay arises from a decaying remnant signal in recently active neurons (Buzsáki, 1989; Foster and Wilson, 2006; Csicsvari et al., 2007; O'Neill et al., 2008). However, it was recently shown that backward replay is not restricted to beginning at the animal's location in the environment (Davidson et al., 2009), which precludes backward replay reflecting only the immediate past. Nevertheless, the environment in Davidson et al. (2009) was a linear track, which means that these replays still occurred over recent trajectories. In fact, all previous reports of backward sequence replay (Foster and Wilson, 2006; Davidson et al., 2009) have been taken from linear tracks in which animals run back and forth, which means that the positions being replayed still reflected recent trajectories taken by the animal. As we will show, backward replay could also reflect remote paths (trajectories not experienced for greater than 10 min), and occurred in large proportions (49% backward to 51% forward) on tracks in which the vast majority of experience occurred in only one direction of travel. In addition, we will show examples of backward replay in which the animal only traveled that complete path in the forward direction, implying that the backward sequential activation of place cells can reflect trajectories never actually experienced by the animal.

One of the key ideas underlying the concept of a cognitive map is that it should enable the planning of novel trajectories never experienced by the animal (Tolman et al., 1946; O'Keefe and Nadel, 1978). Rats can take correctly directed paths involving trajectories that have never been experienced (Tolman et al.,

1946; Alvernhe et al., 2008; D.B. Matthews et al., 1995, Soc. Neurosci., abstract). The cognitive map has been hypothesized to entail a representation of the structure of the environment which enables rats to mentally traverse alternate paths, potentially including trajectories never actually experienced (Tolman, 1948; O'Keefe and Nadel, 1978; Gallistel, 1990; Samsonovich and Ascoli, 2005). Previous studies have identified sequential representations that may reflect such planning (Wood et al., 2000; Ferbinteanu and Shapiro, 2003; Johnson and Redish, 2007; Diba and Buzsáki, 2007; Ainge et al., 2007; Karlsson and Frank, 2009), but all of these previous studies have only identified sequences reflecting trajectories that had been previously experienced. Identification of sequences never experienced by the animal would provide very strong support for the concept of the cognitive map.

Finally, an open question is: how might a map be constructed and maintained to be a balanced representation of the entire environment, even when only parts of the environment are recently visited? Oversampling can produce catastrophic interference and degradation of nonrecently experienced representations. The hippocampus has been suggested to be involved in the prevention of this sort of catastrophic interference (O'Reilly and McClelland, 1994; McClelland et al., 1995). A mechanism by which trajectories that have not been recently visited are preferentially activated in the network would provide a potential solution to this problem.

Previous studies investigating the content of replay have primarily utilized simple tasks in which all trajectories in the environment are equally well experienced and recent. Here we use a new task with multiple possible paths and different behavioral contingencies which manipulates the frequency and recency of experiences to address the relationship between the animal's actual behavior during a given session and the content of what is replayed during the awake state.

RESULTS

To investigate the relationship between experience and the content of hippocampal replay, we recorded neural ensembles from the CA1 region of hippocampus in six rats trained on a spatial decision task with two possible paths (Figure 1A), allowing for the separation of recent and nonrecent experiences. The

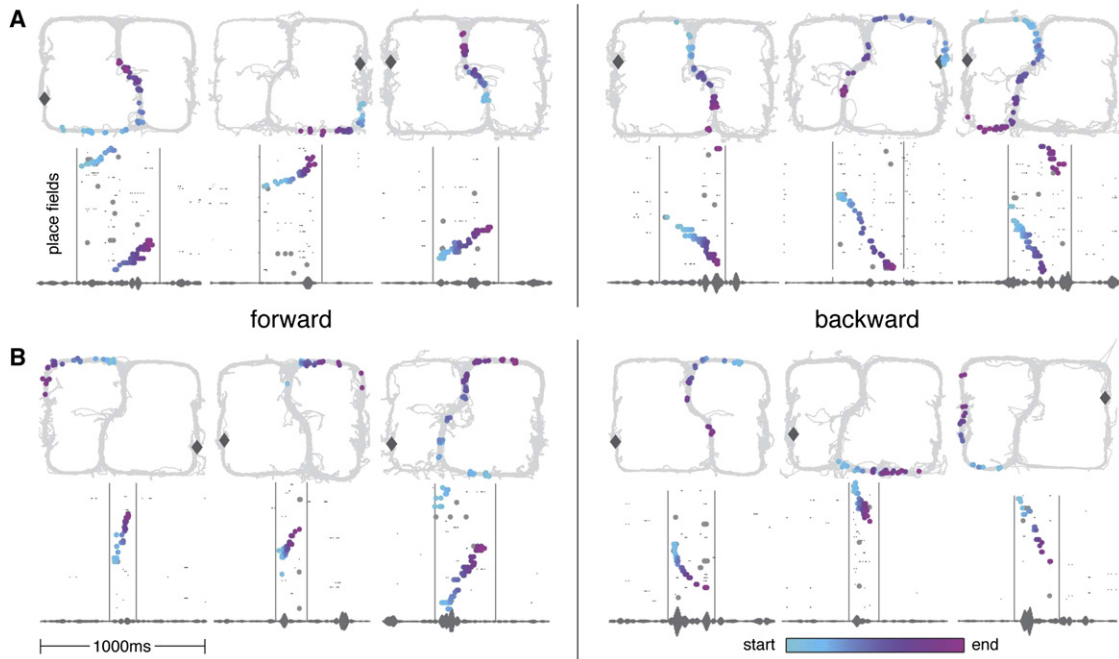


Figure 2. Examples of Forward and Backward Replays

(A) Examples of same-side and central-stem replays (3 fwd, 3 bwd).

(B) Examples of opposite-side replays (3 fwd, 3 bwd).

Gray diamonds indicate the rat's location at the time of the replay. On the bottom panel of each subfigure, spikes are plotted by ordered place field (spatial firing field) center (along either a left or right loop of the maze) over a 1 s period (see [Experimental Procedures](#)). LFPs filtered between 180 Hz and 220 Hz are plotted at the bottom of the panel. Colored points indicate spikes that contribute positively to the sequence score of the replay according to the automatic sequence detection algorithm. The color of the spike indicates its relative time within the replay (light blue, early; light purple, late). Gray points are spikes that do not contribute positively to the score. For cells with multiple place fields, small black points are plotted at every place field center belonging to the cell (colored points occupy the place field center that contributes maximally to the score). Each colored point from the bottom panel is plotted on the 2D maze in the top panel at the location of its 2D place field center. Note that forward replays could begin near the rat's location (A, left) or on the opposite side of the maze (B, left). Similarly, backward replays could begin near the rat's location (A, right) or on the opposite side of the maze (B, right). Backward replays occurred over parts of the environment that were rarely or never experienced in the reverse direction ([Figures S2 and S3](#)).

maze consisted of two choices with a daily changing contingency: the rat was rewarded for turning left (L), turning right (R), or alternating left and right turns (A) at the final choice point. During training sessions, the reward contingency was stable; during recording sessions, it switched approximately midway through the session. Thus, results were analyzed in half-sessions which could be L, R, or A. This means that within a session, rats sampled both loops of the maze, but how long ago each loop was experienced and the total amount of experience on each loop of the maze could differ depending on the contingency. The animals learned the contingency quickly at the onset of the recording session and again after the contingency switch (Figures 1B and 1C). Place cells did not "re-map" after the switch (i.e., place cells maintained the same place fields after the switch; see [Figure S1](#) available online). The vast majority of experience on this track occurred in one direction ([Figures S2 and S3](#)).

Forward and Backward Replay of Local and Nonlocal Trajectories

We analyzed sequences (replays) of place cell activity during SWR events as rats paused at the reward locations. A sequence detection algorithm was applied to quantify the degree to which

spiking activity during SWRs represented a coherent sequence, identifying 1719 replays across 31 sessions from 6 rats (see [Experimental Procedures](#)). The median number of active cells during these events was 10, with a median "sequence score" of 90 (see [Experimental Procedures](#)). In general, replays varied in their starting location and direction and were identified as occurring on either the left, right, or central stem of the maze. In line with previous reports ([Foster and Wilson, 2006](#); [Diba and Buzsáki, 2007](#); [Karlsson and Frank, 2009](#)), we observed replays occurring over spatially proximal portions of the environment (rat sitting on the same side of the maze as the replayed trajectory, "same-side replay," $n = 869$, 450 forward [fwd], 419 backward [bwd]; [Figure 2A](#)). However, we also observed frequent replays of spatially distal portions of the environment (rat on the opposite side of the maze as the replay, "opposite-side replay," $n = 534$, 307 fwd, 227 bwd; [Figure 2B](#)). The median number of active cells (MAC) and the median sequence score (MSS) were similar across all categories (same-side, fwd: MAC = 8 and MSS = 65; same-side, bwd: MAC = 10 and MSS = 85; opposite-side, fwd: MAC = 10 and MSS = 86; opposite-side, bwd: MAC = 11 and MSS = 96). These data show that forward and backward replay robustly occurred over both loops in the

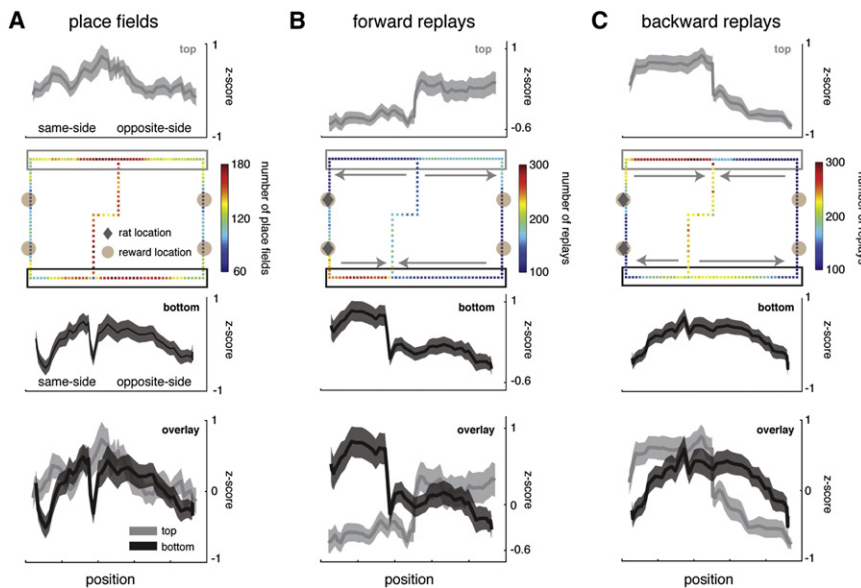


Figure 3. Spatial Distribution of Forward and Backward Replays

(A) Place field spatial distribution.

(B) Forward replay spatial distribution.

(C) Backward replay spatial distribution.

Spatial distributions over the entire environment are shown in the second row. The first and third rows display the spatial distribution over the top and bottom of the maze (indicated by gray and black boxes in row 2), respectively. Spatial distributions over the top and bottom of the maze are overlaid for comparison in the fourth row. All mazes were flipped and aligned such that the animal's location was always on the left side of the maze (indicated by gray diamonds) at the time of the replay. Therefore, the spatial distribution on the left side of the maze reflects same-side replays and the distribution on the right reflects opposite-side replays. The pixel color indicates the total number of replays that represented that particular location in the environment. Error bars in rows 1, 3, and 4 are SEMs over 31 sessions. Note that forward (fwd) and backward (bwd) replays preferentially represented certain portions of the maze, which could not be explained by the place field (pfs) distribution. Overall distributions were significantly different ($p < 10^{-71}$, fwd versus pfs; $p < 10^{-176}$, bwd versus pfs; $p < 10^{-197}$, fwd versus bwd; Kolmogorov-Smirnov tests).

maze when the animal paused at a reward site on one side of the maze. This is consistent with recent studies showing reactivation of experiences from previous environments (Jackson et al., 2006; Gelbard-Sagiv et al., 2008), replay of forward sequences from previous environments (Karlsson and Frank, 2009), and backward replay of trajectories that do not begin at the animal's location on a linear track (Davidson et al., 2009). However, unlike backward replay on linear tracks, this result demonstrates that backward replay robustly occurred in an environment experienced in one direction, over trajectories that had never been experienced in the direction of replay. Although rats would sometimes face backward during early experiences, for each rat there were backward trajectories that were replayed but never experienced and the vast majority (>96%) of all experience was in the forward direction. Yet forward and backward replays were observed with similar proportion.

Spatial Distributions of Forward and Backward Replay

Given that replays could start at any location on the maze, we investigated whether certain portions of the maze were preferentially represented by forward and backward replays. Whereas same-side replays occurred as expected, with forward replays ahead of the animal and backward replays behind it, this pattern was reversed for opposite-side replays (Figure 3). Here, forward replays preferentially occurred on the segment leading up to reward sites, and backward replays similarly covered trajectories ending near reward sites. Thus, distant/nonlocal replays not influenced by the animal's current location did not uniformly represent the environment but preferred certain trajectories. A potential concern is that if the distribution of place fields is nonuniform, more replays will be detected along segments of

the track that have more place fields. To test this possibility, we compared the spatial distributions of forward and backward replay with the spatial distribution of place fields (pfs) and found that they were significantly different ($p < 10^{-71}$, fwd vs. pfs; $p < 10^{-176}$, bwd vs. pfs; Kolmogorov-Smirnov tests). As shown by the place field distribution in Figures 3A and 3B, the trajectory preferences for forward and backward replay on the top and bottom of the maze were not reflected by the spatial distribution of place fields. Backward replays frequently included the central stem, whereas forward replays did so less often. Overall spatial distributions of forward and backward replays were also significantly different ($p < 10^{-197}$; Kolmogorov-Smirnov test). Because trajectories around a given loop were equally experienced, but not equally replayed, this demonstrates a dissociation between the two, which is inconsistent with the hypothesis that the content of replay reflects the amount of previous experience.

Comparison of Data with Current Proposed Mechanisms

Current proposals for the mechanism of replay suggest that it should reflect recent experience (Foster and Wilson, 2006; Jackson et al., 2006) or total experience (Buzsáki, 1989; Wilson and McNaughton, 1994; Redish and Touretzky, 1998; Redish, 1999; Sutherland and McNaughton, 2000). To test these ideas, we examined the distributions of time and number of experiences that had elapsed between replaying a sequence and the most recent experience of that sequence. Because same-side replays involve a recent behavioral experience, we only considered the 534 opposite-side replays for this analysis. Histograms showing the amount of time elapsed since the last experience and the number of laps that had been run since the last experience for each replayed trajectory are shown in Figures 4A




Figure 4. Replay Content Is Incompatible with Scenarios Based on Recent Experience or Accumulated Experience within a Session

(A) Histogram showing the time elapsed since the last behavioral experience over each replayed trajectory.

(B) Histogram showing the number of laps elapsed since the last behavioral experience over each replayed trajectory.

(C) Histogram showing the proportion of total behavioral experience on the same side of the maze as each replayed trajectory at the time of the replay. The actual data (gray bars) are compared against three scenarios, where replay coverage is determined by recent experience (blue line), by accumulated experience (red line), or independent of experience (green line). The curves representing each scenario were constructed based on the animals' actual behavior and replay times (see [Experimental Procedures](#)). Thus, the shapes of the curves reflect the finite session length and the behavioral contingencies. For example, the time-since-last-experience curve for the experience-independent scenario (A) slopes downward due to the decreased probability that a replay event and the last experience over the replayed trajectory will be separated by 2000 s in a 2400 s recording session. Similarly, in (C), the experience-independent curve has three peaks, reflecting the fact that experience was evenly distributed between left and right laps (during alternation) or was primarily on one side of the maze (during left- or right-only contingencies). Note that in (C), only the experience-independent scenario yields a good match to these data, because the other two scenarios fail to account for the peak in replays on the poorly experienced side (black arrow). These results were replicated using a Bayesian decoding approach ([Figure S4](#)).

and 4B. The median time since the last experience was 337 s and the median number of laps since the last experience was 9 laps (366 s, 11 laps for forward replays and 298 s, 7 laps for backward replays). A substantial number of forward and backward trajectories were replayed more than 10 min or 15 laps after they were last experienced.

We also asked whether replayed trajectories favored the side of the maze that had been the most experienced within a session. For each non-central-stem replay (1403 replays), the total number of laps the rat had run on the same side of the maze as the replayed trajectory was divided by the total number of laps that the animal had run on either side of the maze at the time of replay. A histogram of this proportion (the proportion of the rat's total experience that had occurred over the replayed trajectory) is shown in [Figure 4C](#). The symmetry of this histogram indicates that poorly and extensively experienced portions of the maze were replayed with a similar frequency (median of distribution = 0.5 for all replays). Contrary to the predictions of current proposals, replay did not preferentially represent trajectories that had been experienced most often within a session. This also held true for forward replays only and for backward replays only (not shown).

To directly test the recent experience and total experience hypotheses, we compared the data in [Figures 4A–4C](#) with the distributions that would be expected under three scenarios of replay content generation: [1] (blue) replays preferentially include the most recent experience, [2] (red) replays preferentially include trajectories that are highly experienced within a session, and [3] (green) replays are independent of experience (all portions of the environment have an equal chance of being replayed, independent of recent or total experience). For each scenario, we used the actual rat behavior and replay times but reassigned the content of each replay depending on the scenario (see [Experimental Procedures](#)). For example, under scenario 1 (most recent experience), replay was assumed to reflect the most recently experienced lap. Thus, the number of laps since the last experience ([Figure 4B](#)) was always 1, but the time since the last experience ([Figure 4A](#)) depended on the time the animal had been paused at the reward sites.

Our observations are not compatible with the hypothesis that replay reflects the most recent experience, nor with the hypothesis that it reflects the accumulated experience within a session. Instead, our data are most similar to what would be expected if replays were independent of experience within a session. (Data vs. accumulated experience: laps [$p < 10^{-12}$], time [$p < 10^{-10}$], proportion [$p < 10^{-32}$]; data vs. most recent experience: laps [$p < 10^{-54}$], time [$p < 10^{-31}$], proportion [$p < 10^{-34}$]; data vs. experience independent: laps [$p < 0.05$], time [$p < 0.05$], proportion [$p < 0.05$]; Kolmogorov-Smirnov tests; see [Experimental Procedures](#).) Because the number of left and right laps that each animal experienced over its lifetime was counterbalanced, a scenario in which replay equally reflects all experiences over the animal's lifetime would look like the experience-independent scenario. These results are not compatible with the recent or accumulated experience within a session scenarios, but may be explained by a scenario in which replay reflects the animal's total experience over its lifetime. However, the preferential replay of certain trajectories (described above) and the pattern of replays (below) cannot be explained by the experience-independent scenario either.

Task Dependence of the Content of Replay

Because the task contained alternation (A) as well as left-only (L) and right-only (R) half-sessions, we asked whether the proportion of opposite-side replays was different when the animals were alternating compared to when they were performing laps on only one side of the maze. During alternation half-sessions, experiences on the opposite side of the maze occurred more recently in the past and would occur sooner in the future compared with opposite-side experiences during left-only or right-only half-sessions. The median time since the last experience over the replayed trajectory and the number of laps since the last experience for opposite-side replays during alternation was 110 s and 1 lap, whereas the median time and laps since the last experience for opposite-side replays during left-only or right-only half-sessions was 382 s and 11 laps. These distributions were significantly different (time since last experience, alternation vs. left- or right-only: $p < 10^{-4}$; laps since last

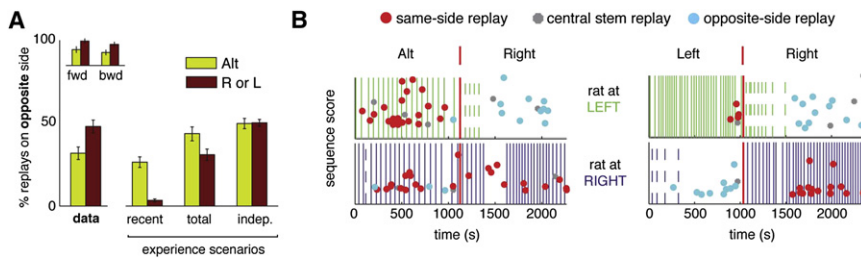


Figure 5. Influence of the Behavioral Contingency on the Content of Replay

(A) Proportion of opposite-side replays on alternation (A) versus right- (R) or left-only (L) contingencies. The “data” group is shown with all replays combined, as well as with forward (fwd) and backward (bwd) replays separately. Whereas the scenarios (described in the text) all predicted a higher or equal proportion of opposite-side replays during alternation half-sessions, the data contained a significantly higher proportion of

opposite-side replays during right- or left-only half-sessions. Error bars are SEMs over half-sessions. This result was replicated using a Bayesian decoding approach (Figure S4).

(B) Examples of same-side and opposite-side replays on L and R contingencies versus an A contingency (two sessions shown). For each session, the top panel shows left lap experiences (vertical green lines), which were followed by pauses at left reward locations, and replays (points) that occurred over the left side of the maze while the animal was paused at either left or right reward locations. The bottom panel shows right lap experiences (vertical purple lines) and replays (points) that occurred over the right side of the maze. Dashed lines indicate error laps (e.g., the animal performs a right lap when only left laps are rewarded). Opposite-side replays are indicated by blue points, same-side replays by red points, and central-stem replays by small gray points. Vertical red lines mark the contingency switch. x axis, time; y axis, replay sequence score (see Experimental Procedures). Note the large number of blue points (opposite-side replays) during L or R contingencies compared to A. Thus, there was a bias for replaying nonlocal trajectories when they were infrequently experienced. This is in contrast to all three scenarios for replay content generation, including the experience-independent scenario.

experience, alternation vs. left- or right-only: $p < 10^{-7}$; Mann-Whitney tests). The recent experience and total accumulated experience within a session scenarios predict an increased number of opposite-side replays during alternation sessions, and the experience-independent scenario predicts an equal proportion of opposite-side replays during alternation and left-only and right-only sessions. Contrary to all three of these scenarios, a larger proportion of opposite-side replays occurred during left-only or right-only half-sessions compared with the alternation half-sessions (Figures 5A and 5B; $p < 0.002$; Mann-Whitney test). This result suggests that replays may serve a role in preserving the representation of nonrecent trajectories. The results shown in Figures 4 and 5 were replicated using a Bayesian one-step decoding method with a 20 ms sliding time window (Zhang et al., 1998; see Figure S4 and Experimental Procedures).

Never-Experienced Shortcut Sequences

Nineteen shortcut sequences were observed over eight sessions from three rats. These sequences traversed a straight path on the top (R152: 1 occurrence; R153: 11 occurrences over 3 sessions) or bottom (R153: 1 occurrence; R158: 6 occurrences over 3 sessions) of the maze between the reward sites. Trajectories spanning the bottom of the maze were very infrequently experienced (R153: 4 times out of approximately 1780 laps; R158: 12 times out of approximately 1550 laps), never rewarded, and at the time of recording a minimum of 3 days had passed since the animal had experienced the trajectory. Trajectories spanning the top of the maze were never experienced by one rat (R153) and only experienced six times by another (R152). All experience on the maze, including all training sessions, was analyzed (Figure S3). Therefore, rat R153 constructed at least 11 never-experienced shortcut sequences across the top of the maze over three separate recording sessions (Figure 6A). Each of these novel trajectories began close to the animal’s location on one side of the maze and ended close to the reward location on the opposite side. Some of the trajectories appear to dip down slightly at the choice point. This is likely to be a conse-

quence of the assignment of place field centers and can be explained by the nonuniform sampling of the region around the choice point (lots of sampling below, but not above) and by the observation that the animals usually cut the corner at the choice point, creating a sampling bias that resulted in downwardly displaced place field centers at the choice point.

In order to control for the possibility that back-to-back forward and backward replays aligned to give rise to what appeared to be a shortcut trajectory, we measured the likelihood that the random alignment of forward and backward replays would account for the number of shortcuts actually observed. The chance of this happening was very small ($p < 10^{-8}$; binomial test; see Table 1). Similarly, the possibility that the novel trajectories were actually two replays in which the second replay picked up where the first one left off was also unlikely to account for the shortcut trajectories ($p < 10^{-5}$; binomial test; see Table 1). Given all possible combinations of forward and backward replays, we calculated the distribution of disjoint, nonshortcut, and shortcut sequences that would occur if forward and backward replays were randomly paired. Both the expected distribution (based on a uniform distribution of forward and backward replays) and the bootstrapped distribution (based on the actual distribution of replays) differed from the observed distribution (Figures 6B and 6C). This strongly suggests that the observed sequences were novel sequence constructions.

DISCUSSION

By considering the contents of what is replayed on a task that manipulated the recency and frequency of experiences, our data speak to two main issues: the first concerns the relationship between replay and experience within a session, and the second concerns the content of what is replayed. We found that, within a session, the hippocampus replayed both frequently and infrequently experienced trajectories in both forward and backward orders. Both forward and backward replay robustly occurred over trajectories not experienced in more than 10 min or 15 laps, in an environment with the vast majority of experience in




Figure 6. Construction of Novel Shortcuts

(A) Examples of novel trajectories. In the bottom panels, spikes are plotted by ordered place field center for both left and right loops over the same 0.5 s period. The gray vertical lines mark the beginning and end of the shortcut sequence and capture the exact same period of time on both left and right loop raster plots (as can also be seen in the repeated LFP trace). The temporally color-coded spikes (as described in the Figure 2 caption) are plotted on the 2D maze (top panels) to visualize the shortcut trajectories spanning the top of the maze.

(B) Examples of a shortcut trajectory (black), a disjoint trajectory (beige), and a nonshortcut sequence (light blue).

(C) Expected, bootstrapped, and observed distributions of disjoint, nonshortcut, and shortcut trajectories (see Experimental Procedures). This analysis shows that the observed shortcuts were extremely unlikely to arise from chance alignments of forward and backward replays, supporting the notion that rats can mentally construct spatially coherent but never-experienced paths (see Figure S3 for paths experienced by each rat).

not be accounted for by the distribution of place fields on the task and suggests that certain trajectories were actively replayed. Proposals suggesting that replay is a simple function of experience cannot explain the preferential replay of particular trajectories.

We also found that the content of replay changed depending on the behavioral task at hand: during alternation half-sessions, a lower proportion of replays occurred along the opposite-side loop compared with left-only and right-only half-sessions, even though the opposite-side loop had been traversed more recently in the past and would occur sooner in the future during the alternation half-sessions. This finding is consistent with the prediction that the content of replay should depend on the animal's behavior, but is the opposite pattern to that predicted by the recency and frequency proposals (compare simulations from Redish and Touretzky, 1998, and the discussion in Foster and Wilson, 2006). Whereas previous studies (Wilson and McNaughton, 1994; Jackson et al., 2006; O'Neill et al., 2006; Karlsson and Frank, 2009) have found increasing replay with experience, these studies did not compare tasks with different behavioral requirements in which parts of the environment were experienced more than others. Therefore, the increase in replay with experience seen in these earlier studies may be due to general experience in the environment rather than the experience of particular trajectories.

one direction. Furthermore, during left-only and right-only half-sessions, trajectories along the nonrecent (opposite-side) loop were replayed with a similar frequency to trajectories on the recent (same-side) loop. This observation was in contrast to alternation half-sessions in which opposite-side loops were replayed less frequently. These observations indicate that current proposals for potential mechanisms of replay that rely on recency or frequency of experience are inadequate.

In theories of learning and consolidation, the information replayed is the information learned (Marr, 1971; Buzsáki, 1989; Wilson and McNaughton, 1994; Redish and Touretzky, 1998; Redish, 1999; Sutherland and McNaughton, 2000; Foster and Wilson, 2006). In our data, trajectories in the environment leading toward reward locations and/or locations where the animal paused after a lap were preferentially replayed. This result could

Previous interpretations of the role of replay have arisen from the concept that it provides a mechanism by which recent experiences are written out from hippocampus to cortex (Marr, 1971; Squire et al., 1984; Squire, 1987; Buzsáki, 1989; Cohen and Eichenbaum, 1993; Alvarez and Squire, 1994; Hasselmo and Bower, 1993; Redish and Touretzky, 1998). Correlational studies have found interactions between hippocampal replay and

Table 1. Comparison of Shortcuts Observed with Expected Proportions of Shortcuts Given Random Pairings of Forward and Backward Replays

Sequence Type	Expected Probability	Bootstrap Probability	Observed Probability	Significance of O-E Difference (Binomial Test)
Disjoint sequences	20/32 = 62.5%	64%	0/19 = 0%	8×10^{-9}
Nonshortcut sequences	8/32 = 25%	24%	3/19 = 15%	0.15
Shortcut sequences	4/32 = 12.5%	12%	16/19 = 84%	2×10^{-12}
Shortcut sequences under back-to-back replay theory	4/12 = 33.3%	33%	16/19 = 84%	7×10^{-6}

E, expected; O, observed.

cortical learning (Hoffman and McNaughton, 2002; Euston et al., 2007; Ji and Wilson, 2007) as well as transfers of dependence (Squire, 1987; Maviel et al., 2004) and learning changes from disruption of SWRs (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010; J.C. Jackson et al., 2009, Soc. Neurosci., abstract). Although our data do not preclude the possibility that information is transferred between hippocampus and cortex during SWRs, it suggests that the information available for transfer is more reflective of the entire set of navigationally available paths rather than the specific experiences themselves.

The results presented above suggest that sequence generation during SWRs is likely to be involved in learning and maintaining a representation of the environment. The forward and backward replay of all trajectories in the environment (not just the recent and well-experienced ones) could be a mechanism to establish and reinforce connections between nearby locations in the representation. During behavior, the combination of phase precession and spike-timing-dependent plasticity has been proposed to enable the storage of sequences in the order experienced by the animal (asymmetric connections are learned) (Levy, 1996; Skaggs et al., 1996; Mehta et al., 1997). Although this mechanism would produce representations of the forward paths traversed by the animal, it would not produce backward connections, nor would it produce novel (never-experienced) connections. Our finding of backward replay in an environment experienced primarily in one direction, over trajectories never experienced in the direction of replay, suggests that backward replay could be an important mechanism for learning a navigationally complete representation of the environment: a representation that reflects not only the trajectories experienced by the animal but also the reverse trajectories available in the environment. This interpretation is consistent with the observation that reactivation and backward replay events are more frequent in novel environments (Foster and Wilson, 2006; Cheng and Frank, 2008), which could be a mechanism to rapidly acquire a complete representation.

Our finding that novel sequences never experienced by the animal are also played out during awake SWRs may also reflect mechanisms for learning the navigationally complete representation of the environment through the internal exploration of potential shortcuts (Samsonovich and Ascoli, 2005). Although they were rare, the likelihood that these novel sequences could be accounted for by chance alignments of forward and backward replays was very low—the distributions observed were significantly different from those expected by a bootstrap pairing of replays. The presence of shortcut sequences offers strong

support that the hippocampal network contains a navigationally complete representation (a cognitive map) of the environment (Tolman, 1948; O'Keefe and Nadel, 1978). These sequences further support the idea that backward replay is not simply an experience replayed in the reverse order, but rather reflects a constructed sequence representing a trajectory in the environment. The properties of sequence “play” described here suggest that hippocampal reactivation during SWRs could be an important part of the mechanism that learns this cognitive map.

One problem with maintaining a stable representation of the environment is the issue of catastrophic interference: when a network simultaneously learns and encodes many sequences, sequences that are not being rehearsed tend to degrade as other activated sequences interfere with them (O'Reilly and McClelland, 1994; McClelland et al., 1995). Preferentially replaying nonlocal trajectories when they are not being visited (as was the case for left-only and right-only half-sessions but not for alternation half-sessions) could be a mechanism to prevent the representation of the nonrecently experienced trajectories from degrading (Pomerleau, 1991). Reactivation of memories from previous tasks within a day (Jackson et al., 2006; Karlsson and Frank, 2009) may serve a similar purpose.

The hippocampal formation has been shown to be important for the imagination of novel situations (Hassabis et al., 2007), and has been implicated in self-projection (Gelbard-Sagiv et al., 2008), the ability to consciously explore the world from different perspectives (Buckner and Carroll, 2007). Our data that sequential place cell activation during awake states reflects forward, backward, and novel sequences spanning the environment, with a pattern more consistent with maintaining a representation of the environment than with replaying recent experiences, support the hypothesis that the hippocampus may provide a potential substrate for self-projection-like processes. The fact that these properties occur during the awake state, while the animal is paused but still engaged in the behavioral task, allows the intriguing speculation that “replay” could contribute to the animal's real-time representation of the world, providing access to information spanning the cognitive map, thereby supporting flexible and goal-driven behaviors.

EXPERIMENTAL PROCEDURES

Subjects

Six male Fisher-Brown-Norway hybrid rats (Harlan; age 7–10 months at time of implantation) were maintained on a synchronous day/night cycle. Rats were

food deprived to no less than 80% of their body weight during behavioral training, and water was freely available in the home cage at all times. All procedures were in accordance with National Institutes of Health guidelines for animal care and were approved by the Institutional Animal Care and Use Committee at the University of Minnesota.

Surgery, Recording, and Histology

After pretraining to proficiency (19–24 days), three rats were implanted with a single-bundle 12-tetrode, two-reference microdrive (Neuro-Hyperdrive; Kopf) directed toward the CA1 hippocampal (HC) subfield (3.8 mm posterior and 2.5 mm right-lateral from bregma). Another three rats were implanted with a double-bundle 12-tetrode, two-reference microdrive directed toward CA1 and ventral striatum (CA1 targets 3.8 mm posterior and 2.5 mm right-lateral from bregma; only HC data were analyzed here). The remaining details of the surgery were as presented previously (Jackson et al., 2006). Tetrodes and references were slowly advanced toward the pyramidal cell layer over approximately 2 weeks after surgery. For HC-only implants, one reference was lowered to the HC fissure and one was left in corpus callosum or a quiet region of cortex to be used as a superficial reference. For the dual-structure implants, the HC reference was placed in the fissure and the ventral striatum reference was placed near corpus callosum; spike data were referenced against the HC reference and local field potential (LFP) data against the ventral striatum reference. Neural activity was recorded and spikes were sorted into putative cells as presented previously (Jackson et al., 2006).

After task performance, rats were overdosed on Nembutal (150 mg/kg, i.p.) and perfused intracardially with formalin. After 24 hr in formalin, brains were transferred to a 30% sucrose-formalin solution, sliced, and stained with cresyl violet using standard procedures. All HC recording locations were verified to lie in the CA1 region of dorsal HC.

The 2T Task

The task consisted of two T intersections, with return rails after the second turn, making it a lap-based task in which the environment was experienced in one direction (Figure 1). Food was delivered at two sites on each return rail contingent on the animal's choice at the second turn. On left laps, the first reward site delivered banana-flavored pellets (5TUL-banana; TestDiet) and on right laps the first reward site delivered fruit-flavored pellets (5TUL-fruit). On both left and right laps, the second reward site delivered white unflavored pellets (5TUL). Each training and recording session lasted 40 min. Training on the task was performed in two phases. During phase 1, rats were trained to run laps on one side of the maze, while the other side was blocked. After running at least 40 laps on 2 consecutive days, phase 2 of training began. Blocks were removed and, on a given day, rats had to run all left laps (L), all right laps (R), or alternating left and right laps (A) in order to receive reward. After consistently getting 80% of the laps correct on all three tasks, rats were implanted with hyperdrives. During initial training, rats attempting to run backward on the maze were blocked by the experimenter. Although there were occasions in which rats faced backward during their early experiences, rats rarely ran backward on the track (Figures S2 and S3). After recovery from surgery, phase 2 of training resumed until the rats regained proficiency and tetrodes were in the cell layer. At this point, training ended and a 6 day sequence of recording sessions began. During the six recording sessions, the task contingency changed approximately midway through the session (mean: 18.07 min \pm 1.13 min; SD). There were six recording sessions to allow for all possible pairings of the three tasks (L-R, R-L, L-A, A-L, R-A, A-R). The order that each rat experienced the six potential pairings was randomized across animals.

Place Fields and Sequences

Cells that fired more than 15,000 spikes over the 40 min session (6.25 Hz) were excluded in the analysis to filter putative inhibitory interneurons. Additionally, cells with fewer than 100 spikes in a session were excluded. Sessions with >40 cells were considered for analysis (31/36 recording sessions), yielding a total of 2183 place cells. For each session separately, position along the maze was linearized separately for left and right laps such that the rat's position along a lap could be described by a single scalar value (Schmitzer-Torbert and Redish, 2004). Place fields were then identified as contiguous linear pixels

(one linear pixel is approximately 3.5 cm along the linearized maze) with average activity > 5% of the maximum rate observed over the session for any cell at any pixel (cells could have more than one place field, although place fields separated by only a single pixel were merged). In this way, we identified 3088 place fields from the 2183 place cells. This allowed for the determination of place field centers along left and right laps. The centers were ordered from maze start (MS) to MS in the direction the rats traveled around the maze. The following algorithm was then performed four times separately for each session (left lap centers, forward [fwd] replays; left lap centers, backward [bwd] replays; right lap centers, fwd replays; right lap centers, bwd replays). Replays had to pass two sets of criteria before being included: first detection and then significance testing.

Sequences were identified using an algorithm that detects sequence structure in the pattern of place cell activity by comparing the times and place cell centers of spike pairs occurring in a flexible time window (see Supplemental Experimental Procedures). This algorithm resulted in a series of time windows (start and stop times for each replay), place field center-labeled spikes, and scores for each forward and backward sequence occurring on left and right laps during a single session. These sequences were then analyzed to identify significant sequence replays (using two independent bootstrapping procedures) occurring during SWR complexes, while animals were paused at reward sites (see Supplemental Experimental Procedures).

The spatial coverage of each replay was determined from the active place fields contributing positively to the sequence score and was described by a vector of ones and zeros, with each element representing a pixel in the environment. Replays were identified as occurring on one side of the maze (>10% coverage on the non-central stem of one side and <5% coverage on the other) or else were labeled as central-stem replays. If a replay occurred over the right portion of the maze while the rat was at a right-side reward site, it was labeled a same-side replay. If the rat was at a right-side reward site and the replay covered a trajectory on the left side of the maze, it was labeled an opposite-side replay. Same-side and opposite-side replays were defined analogously when the rat was sitting at a reward site on the left side of the maze. The vector representation of each replay was used to create the spatial distribution plots in Figure 3.

Scenarios

Three scenarios based on three proposals ([1] replays reflect the most recent experience; [2] replays reflect accumulated experience within a session; [3] replays are independent of experience) were compared with our data. As described above, each observed replay was classified as occurring on the left, right, or central-stem portion of the maze. This information was then used to identify same-side and opposite-side replays based on the rat's location at the time of the replay. In order to generate a comparable data set that was representative of each scenario, the actual replay times and rat locations were kept constant and the left/right classification was changed for each individual replay (central-stem replays were left intact). For scenario 1, replays were restricted to cover only the parts of the maze that were traversed over the last lap (i.e., all replays occurring while the animal is paused at a reward location on the left side of the maze after performing a left lap were assigned to the left side of the maze). For scenario 2, replays stochastically preferred portions of the maze that were most visited at the time of the replay (i.e., regardless of the rat's paused location on the maze, the probability of a replay being assigned to the left side of the maze was equal to the proportion of accumulated experience on the left side of the maze at the time of the replay). For scenario 3, each non-central-stem replay was randomly assigned to be on the left or right, independent of recent or accumulated experience during the session. From these new assignments, same-side and opposite-side classifications were made and time since last experience, laps since last experience, and proportion of accumulated experience on the side of the replay were calculated for each scenario. Each of these distributions was calculated 100 times and compared to the actual data using Kolmogorov-Smirnov tests. The mean p value and mean distribution are reported for each comparison (see Figures 4A–4C and 5A).

Additionally, we performed one-step Bayesian decoding (Zhang et al., 1998) with a uniform spatial prior and a 20 ms sliding time window on the identified replay times. Replays were classified as occurring on the side of the maze

with the largest cumulative probability, and the results from Figures 4A–4C and 5A were replicated (Figure S4). See Supplemental Experimental Procedures for a discussion of our decoding method compared with Bayesian decoding.

Shortcuts

A shortcut sequence on the two-choice maze can be viewed as a combination of one forward and one backward replay that are temporally and spatially adjacent on the top or bottom of the maze. There are 32 possible combinations of fwd and bwd replays on the maze. There are four types of fwd replays: top left (TL), top right (TR), bottom left (BL), and bottom right (BR). Each of the fwd replays can be paired with the four types of bwd replay (TL, TR, BL, BR), to make $4 \times 4 = 16$ combinations. Because order matters, there are $16 \times 2 = 32$ possible combinations. Out of these combinations, 4 give rise to shortcut sequences (bwdTR-fwdTL, bwdTL-fwdTR, fwdBR-bwdBL, fwdBL-bwdBR), 8 give rise to a coherent trajectory that does not result in a shortcut path, and 20 do not create a coherent trajectory at all (see Figure 6). Thus, the theoretical probabilities of shortcut, nonshortcut, and disjoint combinations are 4/32, 8/32, and 20/32, respectively. These probabilities assume that fwd and bwd replays are evenly distributed on the top and bottom of the maze. We used a bootstrapping procedure to determine the actual probabilities of each combination from the distributions of replays observed. This procedure randomly paired each fwd replay with 100 bwd replays and each bwd replay with 100 fwd replays. The probability of shortcut, nonshortcut, and disjoint combinations were then calculated from the number of occurrences of the 32 possible combinations.

A simple algorithm was then used to identify the actual numbers of shortcut, nonshortcut, and disjoint combinations present in the data. The algorithm first identified replays that occurred primarily in each corner of the maze (TL, TR, BL, BR). Using this set of replays, the algorithm identified all pairs of fwd and bwd replays with midtimes separated by <150 ms (the mean duration of the replays). Other methods of identification produced similar results. The algorithm identified 16 shortcut sequences, 3 nonshortcut sequences, and 0 disjoint pairings. To test the likelihood that 16/19 pairings would be contiguous sequences based on the replay distributions in our data, we calculated the probability that the random alignment of fwd and bwd replays could give rise to 0/19 disjoint sequences (binomial test; see Table 1). Out of the 32 combinations, 12 give rise to sequences (back-to-back replays in which the second replay picks up where the first left off), 4 of which are shortcut sequences. To test whether the shortcut sequences resulted from two back-to-back replays in which the second replay picked up where the first one left off, we calculated the probability under this theory that 16/19 sequence pairings would result in shortcuts (binomial test; see Table 1).

SUPPLEMENTAL INFORMATION

Supplemental Information includes four figures, four movies, and Supplemental Experimental Procedures and can be found with this article online at doi:10.1016/j.neuron.2010.01.034.

ACKNOWLEDGMENTS

We thank Chris Boldt and Adam Steiner for technical assistance, Adam Johnson for comments on a previous version of the manuscript, and the members of the Redish lab for discussion. This work was funded by NIH grant R01 MH-080381, the Pennsylvania Department of Health, and NSF IGERT DGE-0549352.

Accepted: January 8, 2010

Published: March 10, 2010

REFERENCES

Ainge, J.A., Tamosiunaite, M., Woergoetter, F., and Dudchenko, P.A. (2007). Hippocampal CA1 place cells encode intended destination on a maze with multiple choice points. *J. Neurosci.* 27, 9769–9779.

Alvarez, P., and Squire, L.R. (1994). Memory consolidation and the medial temporal lobe: a simple network model. *Proc. Natl. Acad. Sci. USA* 91, 7041–7045.

Alvernhe, A., Van Cauter, T., Save, E., and Poucet, B. (2008). Different CA1 and CA3 representations of novel routes in a shortcut situation. *J. Neurosci.* 28, 7324–7333.

Buckner, R.L., and Carroll, D.C. (2007). Self-projection and the brain. *Trends Cogn. Sci.* 11, 49–57.

Buzsáki, G. (1989). Two-stage model of memory trace formation: a role for “noisy” brain states. *Neuroscience* 31, 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.

Cheng, S., and Frank, L.M. (2008). New experiences enhance coordinated neural activity in the hippocampus. *Neuron* 57, 303–313.

Cohen, N.J., and Eichenbaum, H. (1993). *Memory, Amnesia, and the Hippocampal System* (Cambridge, MA: MIT Press).

Csicsvari, J., O’Neill, J., Allen, K., and Senior, T. (2007). Place-selective firing contributes to the reverse-order reactivation of CA1 pyramidal cells during sharp waves in open-field exploration. *Eur. J. Neurosci.* 26, 704–716.

Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). Hippocampal replay of extended experience. *Neuron* 63, 497–507.

Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nat. Neurosci.* 10, 1241–1242.

Ego-Stengel, V., and Wilson, M.A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus* 20, 1–10.

Euston, D.R., Tatsuno, M., and McNaughton, B.L. (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science* 318, 1147–1150.

Ferbinteanu, J., and Shapiro, M.L. (2003). Prospective and retrospective memory coding in the hippocampus. *Neuron* 40, 1227–1239.

Foster, D.J., and Wilson, M.A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* 440, 680–683.

Gallistel, C.R. (1990). *The Organization of Learning* (Cambridge, MA: MIT Press).

Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R., and Fried, I. (2008). Internally generated reactivation of single neurons in human hippocampus during free recall. *Science* 322, 96–101.

Girardeau, G., Benchenane, K., Wiener, S.I., Buzsáki, G., and Zugaro, M.B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nat. Neurosci.* 12, 1222–1223.

Hassabis, D., Kumaran, D., Vann, S.D., and Maguire, E.A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *Proc. Natl. Acad. Sci. USA* 104, 1726–1731.

Hasselmo, M.E., and Bower, J.M. (1993). Acetylcholine and memory. *Trends Neurosci.* 16, 218–222.

Hoffman, K.L., and McNaughton, B.L. (2002). Coordinated reactivation of distributed memory traces in primate neocortex. *Science* 297, 2070–2073.

Jackson, J.C., Johnson, A., and Redish, A.D. (2006). Hippocampal sharp waves and reactivation during awake states depend on repeated sequential experience. *J. Neurosci.* 26, 12415–12426.

Jensen, O., and Lisman, J.E. (2000). Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase encoding. *J. Neurophysiol.* 83, 2602–2609.

Ji, D., and Wilson, M.A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat. Neurosci.* 10, 100–107.

Johnson, A., and Redish, A.D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.* 27, 12176–12189.

Johnson, A., Jackson, J., and Redish, A.D. (2008). Measuring distributed properties of neural representations beyond the decoding of local

- variables—implications for cognition. In *Mechanisms of Information Processing in the Brain: Encoding of Information in Neural Populations and Networks*, C. Hölscher and M.H.J. Munk, eds. (Cambridge, UK: Cambridge University Press), pp. 95–119.
- Karlsson, M.P., and Frank, L.M. (2009). Awake replay of remote experiences in the hippocampus. *Nat. Neurosci.* *12*, 913–918.
- 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.
- Lee, A.K., and Wilson, M.A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* *36*, 1183–1194.
- Levy, W.B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus* *6*, 579–591.
- Marr, D. (1971). Simple memory: a theory of archicortex. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* *262*, 23–81.
- Maviel, T., Durkin, T.P., Menzaghi, F., and Bontempi, B. (2004). Sites of neocortical reorganization critical for remote spatial memory. *Science* *305*, 96–99.
- McClelland, J.L., McNaughton, B.L., and O'Reilly, R.C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* *102*, 419–457.
- 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.
- Nadasdy, Z., Hirase, H., Czurko, A., Csicsvari, J., and Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *J. Neurosci.* *19*, 9497–9507.
- Nakashiba, T., Buhl, D.L., McHugh, T.J., and Tonegawa, S. (2009). Hippocampal CA3 output is crucial for ripple-associated reactivation and consolidation of memory. *Neuron* *62*, 781–787.
- 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'Neill, J., Senior, T., and Csicsvari, J. (2006). Place-selective firing of CA1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. *Neuron* *49*, 143–155.
- O'Neill, J., Senior, T.J., Allen, K., Huxter, J.R., and Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nat. Neurosci.* *11*, 209–215.
- O'Reilly, R.C., and McClelland, J.L. (1994). Hippocampal conjunctive encoding, storage, and recall: avoiding a trade-off. *Hippocampus* *4*, 661–682.
- Pomerleau, D.A. (1991). Efficient training of artificial neural networks for autonomous navigation. *Neural Comput.* *3*, 88–97.
- Redish, A.D. (1999). *Beyond the Cognitive Map: From Place Cells to Episodic Memory* (Cambridge, MA: MIT Press).
- Redish, A.D., and Touretzky, D.S. (1998). The role of the hippocampus in solving the Morris water maze. *Neural Comput.* *10*, 73–111.
- Redish, A.D., McNaughton, B.L., and Barnes, C.A. (1998). Reconciling Barnes et al. (1997) and Tanila et al. (1997a,b). *Hippocampus* *8*, 438–443.
- Samsonovich, A.V., and Ascoli, G.A. (2005). A simple neural network model of the hippocampus suggesting its pathfinding role in episodic memory retrieval. *Learn. Mem.* *12*, 193–208.
- Schmitzer-Torbert, N., and Redish, A.D. (2004). Neuronal activity in the rodent dorsal striatum in sequential navigation: separation of spatial and reward responses on the multiple T task. *J. Neurophysiol.* *91*, 2259–2272.
- 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*, 149–173.
- Squire, L.R. (1987). *Memory and Brain* (New York: Oxford University Press).
- Squire, L.R., Cohen, N.J., and Zola-Morgan, M. (1984). The medial temporal lobe memory system. In *Memory: From Mind to Molecules*, eds. H. Weingartner and E.S. Parker, (Hillsdale, NJ: Erlbaum), pp. 185–210.
- Sutherland, G.R., and McNaughton, B.L. (2000). Memory trace reactivation in hippocampal and neocortical neuronal ensembles. *Curr. Opin. Neurobiol.* *10*, 180–186.
- Tolman, E.C. (1948). Cognitive maps in rats and men. *Psychol. Rev.* *55*, 189–208.
- Tolman, E.C., Ritchie, B.F., and Kalish, D. (1946). Studies in spatial learning. I. Orientation and the short-cut. *J. Exp. Psychol.* *36*, 13–24.
- Tse, D., Langston, R.F., Kakeyama, M., Bethus, I., Spooner, P.A., Wood, E.R., Witter, M.P., and Morris, R.G.M. (2007). Schemas and memory consolidation. *Science* *316*, 76–82.
- 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.
- Ylinen, A., Bragin, A., Nadasdy, Z., Jando, G., Szabo, I., Sik, A., and Buzsáki, G. (1995). Sharp wave-associated high-frequency oscillation (200 Hz) in the intact hippocampus: network and intracellular mechanisms. *J. Neurosci.* *15*, 30–46.
- Zhang, K., Ginzburg, I., McNaughton, B.L., and Sejnowski, T.J. (1998). Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J. Neurophysiol.* *79*, 1017–1044.