

# Hippocampal Replay Captures the Unique Topological Structure of a Novel Environment

Xiaojing Wu and David J. Foster

Solomon H. Snyder Department of Neuroscience, Johns Hopkins University School of Medicine, Baltimore, Maryland 21205

Hippocampal place-cell replay has been proposed as a fundamental mechanism of learning and memory, which might support navigational learning and planning. An important hypothesis of relevance to these proposed functions is that the information encoded in replay should reflect the topological structure of experienced environments; that is, which places in the environment are connected with which others. Here we report several attributes of replay observed in rats exploring a novel forked environment that support the hypothesis. First, we observed that overlapping replays depicting divergent trajectories through the fork recruited the same population of cells with the same firing rates to represent the common portion of the trajectories. Second, replay tended to be directional and to flip the represented direction at the fork. Third, replay-associated sharp-wave-ripple events in the local field potential exhibited substructure that mapped onto the maze topology. Thus, the spatial complexity of our recording environment was accurately captured by replay: the underlying neuronal activities reflected the bifurcating shape, and both directionality and associated ripple structure reflected the segmentation of the maze. Finally, we observed that replays occurred rapidly after small numbers of experiences. Our results suggest that hippocampal replay captures learned information about environmental topology to support a role in navigation.

## Introduction

The hippocampus is critical for learning and memory, particularly spatial memory (Olton and Samuelson, 1976; O'Keefe and Nadel, 1978) and memory for events (Scoville and Milner, 1957; Gaffan, 1994; Steele and Morris, 1999). Neurons in the hippocampus fire at specific locations in the environment, and hence are termed "place cells" (O'Keefe and Dostrovsky, 1971; O'Keefe and Nadel, 1978). It is now well established that hippocampal place cells engage in extended spiking sequences across neuronal ensembles that match sequences from previous behavioral episodes, often in a temporally condensed form during 100–200 ms sharp-wave-ripple (SWR) events in the hippocampal local field potential (LFP; Louie and Wilson, 2001; Lee and Wilson, 2002; Foster and Wilson, 2006; Csicsvari et al., 2007; Diba and Buzsáki, 2007; Ji and Wilson, 2007; Davidson et al., 2009; Karlsson and Frank, 2009; Gupta et al., 2010). While SWR-associated place-cell sequences are associated with memory consolidation, particularly during sleep (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010; O'Neill et al., 2010; Carr et al., 2011), there is growing evidence that they also contribute to navigational learning and planning (Buckner, 2010; Foster and Knierim, 2012; Jadhav et al., 2012). Place-cell sequences occur during the awake state during pauses in behavior at reward sites, in which the previously expe-

rienced behavioral sequence is replayed in reverse order (Foster and Wilson, 2006; Diba and Buzsáki, 2007; Davidson et al., 2009; Karlsson and Frank, 2009), and in a manner likely to be modulated by reward (Singer and Frank, 2009), hence providing an ideal representation for associating locations with predictions of expected future reward (Montague et al., 1996; Foster and Wilson, 2006). More recently, it has been demonstrated that SWR-associated place-cell sequences occurring immediately before movement in a spatial memory task depict the future trajectory that the animal will take to the remembered goal location (Pfeiffer and Foster, 2013). Therefore, hippocampal SWR-associated place-cell sequences might provide a mechanism by which the brain addresses the learning, memory, and planning demands inherent in memory-based navigation (Tolman, 1948; O'Keefe and Nadel, 1978).

A key aspect of the navigation problem is dealing with the topological structure of the terrain, that is, which places are connected to, or accessible from, which other places. This structure determines the set of available paths that can be traversed, as well as the barriers that must be avoided. Navigational schemes that ignore topological structure, such as simple dead reckoning, can work well in open, unobstructed environments such as those encountered by certain species of desert ant (Gallistel, 1990), but in more complex environments successful navigation is likely to require modes of planning that incorporate topological structure (O'Keefe and Nadel, 1978).

We hypothesized that hippocampal SWR-associated place-cell sequences would capture the spatial topology of the environment, rather than capturing only the temporal structure of experiences in the environment as independent episodes. We tested this hypothesis by recording place-cell activity while rats explored a novel environment with an unpredictable structure: a

Received Aug. 9, 2013; revised March 13, 2014; accepted March 20, 2014.

Author contributions: X.W. and D.J.F. designed research; X.W. performed research; X.W. and D.J.F. analyzed data; X.W. and D.J.F. wrote the paper.

This work was supported by the Alfred P. Sloan Foundation, the Whitehall Foundation, and National Institutes of Health Grant R01 MH085823.

The authors declare no competing financial interests.

Correspondence should be addressed to Solomon H. Snyder Department of Neuroscience, Johns Hopkins University School of Medicine, Baltimore, MD 21205. E-mail: david.foster@jhu.edu.

DOI:10.1523/JNEUROSCI.3414-13.2014

Copyright © 2014 the authors 0270-6474/14/346459-11\$15.00/0

Y-shaped maze with asymmetrical arm lengths. Our findings suggest that SWRs and associated place-cell sequences do indeed capture the learned topological structure of the environment.

## Materials and Methods

**Electrophysiology and behavioral task.** Male Long–Evans rats, 3.5–4.5 months in age, were pretrained to run back and forth on a linear track for chocolate milk reward available at both track ends under moderate food deprivation. Pretraining was ended when animals reached the performance criterion of running ~30 laps within 30–40 min, usually after 1–2 weeks. Custom-designed microdrives consisting of 40 independently adjustable tetrodes were then implanted with half the tetrodes targeting the right dorsal CA1 area (−3.6, −2.2 mm from bregma), and the other half targeting the right medial prefrontal cortex (+3.2, −0.8 mm from bregma; data not presented). Tetrodes were gradually moved into the CA1 pyramidal cell layer with one tetrode left in the white matter above the cell layer to provide a reference signal to all other tetrodes.

When all tetrodes were in position, the first recording day, on which rats were first exposed to a modified Y maze, began. The Y maze was composed of one long arm (145 cm run segment) and two short arms (65 cm run segment), all of which were separated by 120°. Each arm had a wider reward area at the end (16 cm in length). One short arm was chosen to be the central (C) arm; the other two arms [right (R); left (L)] were termed “alternating arms.” The rat was placed at the baited end of the central arm, and was allowed to freely explore the Y maze and rewarded according to the following alternation rule: the first arrival at an alternating end was rewarded; thereafter, returns to the central arm were rewarded while visits to the alternating arms were rewarded only if the arm identity was different from the last alternating arm rewarded.

Spiking activities and LFP signals were recorded (Neuralynx) and on-line digitally filtered in different frequency bands (600–6000 Hz for spikes and 0.1–500 Hz for LFP). Rats’ positions were signaled by a red LED and a green LED attached to the drive, and were recorded from an overhead camera. Recording sessions lasted 1–2.5 h and were ended when rats stopped running.

Recordings lasted 6–8 d with one session run on each day and were terminated when cell yield became poor. Rats were killed, and lesions were made on selected tetrodes by passing current (30  $\mu$ A for ~5 s) through each tetrode. Animals were then perfused with formalin. Brains were removed, sectioned, and stained for cresyl violet to verify tetrode recording locations.

The spike data were manually clustered in custom software Xclust2 (written by M. Wilson, Massachusetts Institute of Technology, Cambridge, MA). Each visually discernable cluster was separated by gradually refining its boundary iteratively in each plot of spike waveform peak amplitudes on a pair of tetrode channels (i.e., each of the six two-dimensional projections in the four-dimensional spike peak amplitude space). A cluster was considered as a putative pyramidal single unit when its complex spike index (CSI), a measurement for the bursting tendency of the unit characteristic of hippocampal pyramidal neurons (Nakazawa et al., 2002), was >5%; a cluster was considered to be well isolated and largely free of noise when its cross-CSI with every other cluster from the same tetrode was <3%. Many of the clusters included in our study had single-unit CSIs >10% and cross-CSIs <1%. Cluster stability was verified with the plot of peak amplitudes against time during the entire recording session. Data from the very first sessions of Rats 1 and 3 are presented in this article. Rat 2 did not explore all three arms in his first session (he made only one trip to the reward area of the R arm and ran back to the C reward area, where he stayed during the rest of the session), and data from his second session on the following day are presented here.

**Position linearization and place-field computation.** Recorded positions were projected onto three centerlines (defined by the experimenter) aligned with the three arms of the Y maze. The three lines were then concatenated to produce a linear axis (Fig. 1). Nondirectional place fields were computed by using all spike data (1.8–1.9 cm position bins) and smoothed with a Gaussian filter (SD = 5 bins). Note that firing rates on the three arms were separately smoothed to minimize misestimation of place fields around the choice point. Putative interneurons (mean peak-

to-trough spike width <0.34 ms) and neurons with insignificant place fields (peak firing rate on the Y maze  $\leq$  1 Hz) were excluded from a total of 99, 79, and 62 single units recorded from Rats 1–3. The rest of the single units were considered putative place cells (88 from Rat 1, 67 from Rat 2, 58 from Rat 3) and were used in the following analyses.

**Candidate events and trajectory-specific subregions.** For each recording session, a smoothed spike density function was computed with all spikes from all putative place cells (10 ms time bins; Gaussian filter SD = 15 ms). Candidate events were defined as epochs of spikes during which spike densities were above the mean of the function, and contained peaks above 2 SDs over the mean. Only candidate events that occurred when a rat’s speed was <5 cm/s were considered (2514 from Rat 1, 3222 from Rat 2, 2902 from Rat 3). A Bayesian decoding algorithm (Davidson et al., 2009) was then applied to the candidate events, which calculated the probability of the ensemble of neurons representing each position bin during each time bin.

Single-arm replays representing each individual arm and three types of joint replay extended across each pair of arms (C  $\leftrightarrow$  R, C  $\leftrightarrow$  L, and R  $\leftrightarrow$  L, referred to as CR, CL, and RL) were considered. To maximize the likelihood of detecting all six types of replay in an unbiased manner, we segmented the posterior probability matrix of each candidate event in both position and time to further define trajectory-specific subregions. Each candidate event was first separated along the position axis into three segments corresponding to the three arms (Fig. 1C). We defined a maximum a priori probability (MAP) function as the largest probability across all positions per time bin, smoothed in time with a Gaussian filter (SD, 10 ms; Fig. 1C center, blue curves). Each single-arm segment was then segmented in time, and a trajectory-specific subregion was created around the largest peak of the MAP function to include time bins whose MAP values were above a threshold of five times the chance level, namely  $(1/\text{total \# position bins}) \times 5$  (a fixed threshold only dependent on the size of the track and not the quality of neuronal data). The trajectory-specific subregions of the single arms were then combined in pairs to create those of the joint arms, which expanded from the earliest to the latest time bins of the two corresponding single-arm subregions and contained all position bins belonging to the two corresponding single arms (Fig. 1C, right). It should be noted that all single-arm and joint-arm subregions were continuous in time (i.e., arm segments were not shifted across time bins so that they could be pieced together to form a longer sequence); thus, this segmentation method does not artificially create coherent coding of one arm and then another.

**Replay identification.** To determine whether the posterior probabilities within a trajectory-specific subregion gave rise to a replay sequence, the following three variables were calculated: (1) length of subregion in time; (2) arm coverage [a position bin was considered to be represented in the subregion if its largest probability across time is above the threshold  $(1/\text{total number of position bins}) \times 5$ , and the percentage of represented positions out of all positions within the subregion was defined to be the arm coverage]; and (3) weighted correlation, which is an adapted form of the Pearson’s correlation, measuring the strength of correlation between the changes in probability values across time and position, and using all pixels in the subregion, given by the following:

$$\text{Weighted mean: } m(x; w) = \frac{\sum_{i=1}^M \sum_{j=1}^N w_{ij} x_i}{\sum_{i=1}^M \sum_{j=1}^N w_{ij}}$$

$$\text{Weighted covariance: } \text{cov}(x, y; w) =$$

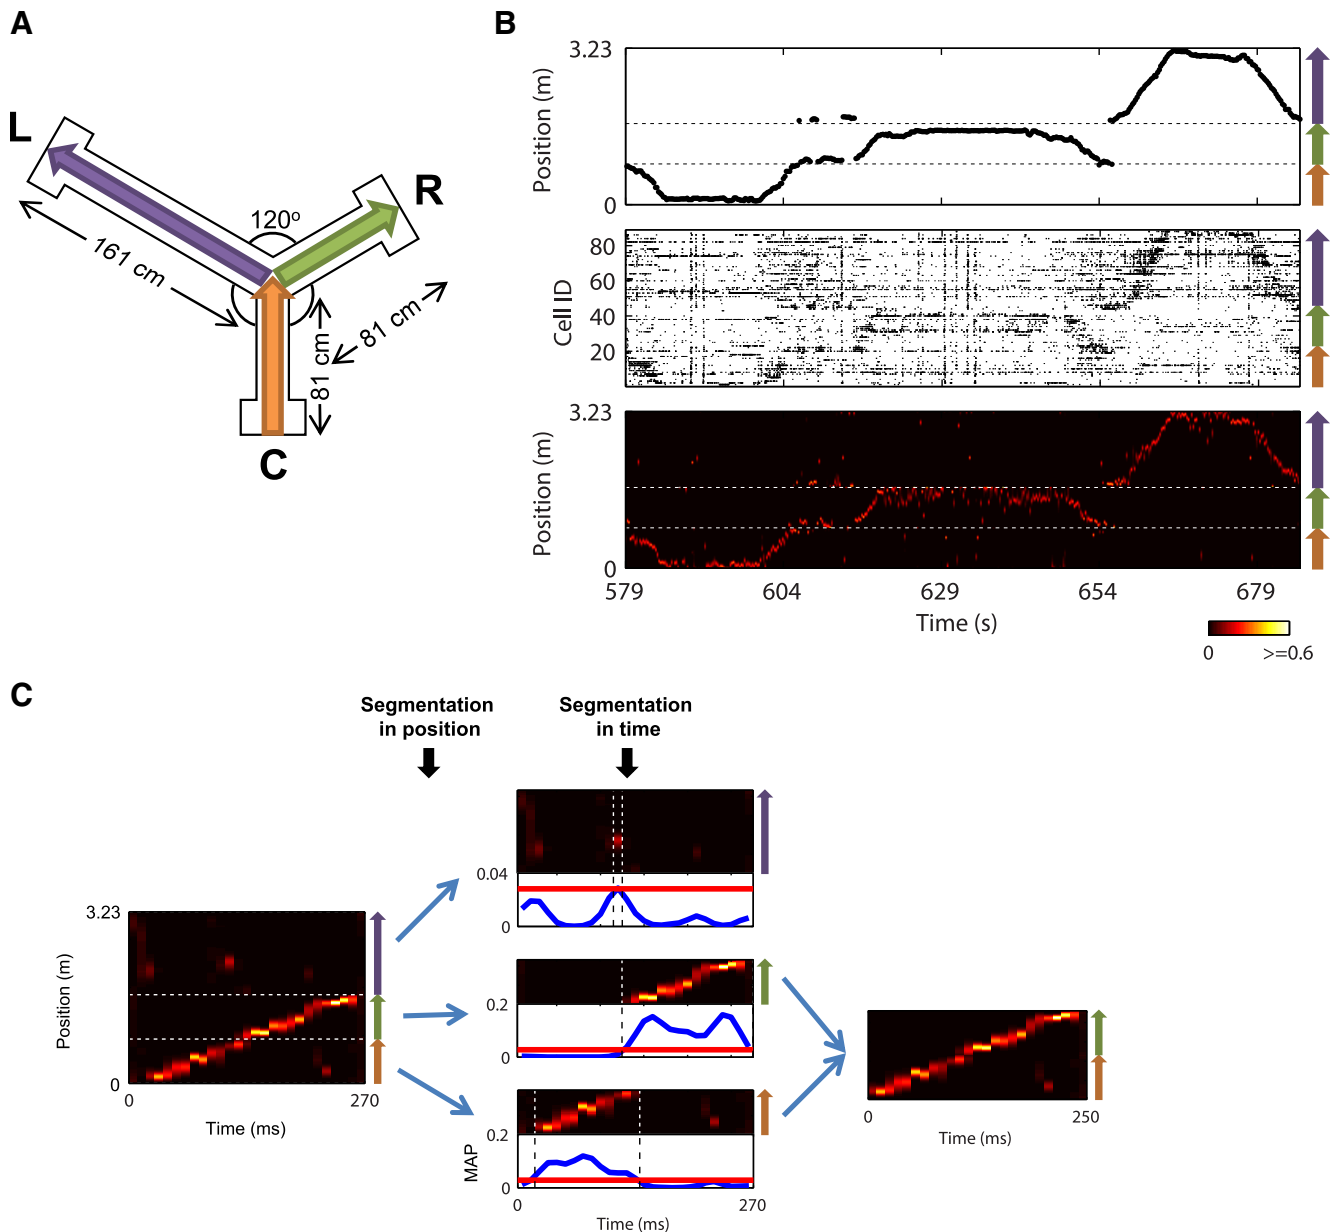
$$\frac{\sum_{i=1}^M \sum_{j=1}^N w_{ij} (x_i - m(x; w))(y_j - m(y; w))}{\sum_{i=1}^M \sum_{j=1}^N w_{ij}}$$

$$\text{Weighted correlation: } \text{corr}(x, y; w) =$$

$$\text{cov}(x, y; w) / \sqrt{\text{cov}(x, x; w)\text{cov}(y, y; w)}$$

Where  $x_i$  is the  $i$ th time bin,  $y_j$  is the  $j$ th position bin,  $w_{ij}$  is the probability of pixel  $(i, j)$ , and  $M$  and  $N$  are the total numbers of the time and position bins of a given subregion.

A single-arm subregion was determined to contain the corresponding single-arm replay if the following three criteria were met: (1) subregion



**Figure 1.** Place-cell activities in modified Y maze. **A**, Modified Y maze. **B** An example epoch of 106 s of recording. Top, Linearized running trajectory of Rat 1. Horizontal dashed lines indicate arm boundaries. Colored arrows on right indicate arm alignment along linear axis (same below). The rat ran through center → C → center → R → center → L → center. Middle, Simultaneously recorded spike trains from 88 putative place cells, ordered by locations of the peak firing rates of the cells on linearized maze. Bottom, Position estimation based on these spikes, where posterior probabilities of position representations in each 250 ms window were indicated by a hot scale. **C**, Illustration of definition of trajectory-specific subregions within a candidate event for replay detection. An example candidate event (left) was first segmented in position into three segments, with each one corresponding to an individual arm: center top, the L arm; center middle, the R arm; and center bottom, the C arm. Blue curves are MAP functions, defined as the largest posterior probability across all positions per time bin, and calculated separately for each segment. Vertical dashed lines indicate the windows within which MAPs were above the threshold  $(1/\text{total number of position bins}) \times 5$ , which is marked by red lines, thus defining trajectory-specific subregions for the three single-arm segments. In this example, only C and R subregions passed all criteria for containing replay structure, which were combined to define the subregion for the joint-arm path CR (right). This candidate event was finally determined to contain a joint replay of CR.

length  $\geq 50$  ms (five time bins); (2) arm coverage of  $>50\%$ ; and (3) absolute weighted correlation  $> 0.5$ .

A joint-arm subregion was determined to contain the corresponding joint replay if the following two criteria were met: (1) absolute weighted correlation  $> 0.5$ ; and (2) both of the constituent single-arm subregions meet criteria 1 and 2 for single-arm replays and have the same signs of weighted correlation as that of the joint-arm subregion.

Finally, for a candidate event to be considered to contain a single-arm replay, none of the joint-arm subregions should contain any replays; for a candidate event to be considered to contain a joint replay, neither of the other joint-arm subregions should contain any replays. It is important to

note that the detection of replay structure within a candidate event is highly dependent on the recorded cell assembly, as the recording of an insufficient number of cells participating in an actual replay event will likely generate a false-negative result.

*Cumulative replay numbers and place-field modular shuffle.* To quantify replay occurrence, we focused on changes in replay number across stopping periods, periods of time spent between runs at a reward area, where the majority of reactivation activities occurred. For each joint replay type, the cumulative number of identified replays was counted for each stopping period, starting from the one immediately following the first lap of traversal across the corresponding arms.

A place-field modular shuffle was used to determine the stopping period by which the cumulative number of a replay first reached a significant level, which circularly shifted each cell's place field by a random number of position bins. This method preserved each cell's spike train and local place-field structure. The same algorithm described in the previous two sections was applied to cells' shuffled place fields and original spike trains to create sample distributions of cumulative replay numbers (5000 shuffles), from which Monte Carlo  $p$  values of the original cumulative numbers were calculated. The first significant stopping period was defined by  $p < 0.05$ . The number of laps (running from one arm end to another) of corresponding arm traversals was used to quantify the amount of physical experience acquired before the first significant stopping period. The total number of replays in each session, from original and shuffled data, was used for Figure 2E.

**Comparison of cell activities in joint replay common segments.** To understand the underlying neuronal firing patterns of joint replays, we compared individual neuronal firing activities during the overlapping segments of joint replays that represented a common arm, between joint replays depicting diverging paths. Take joint CR and CL replays for example. First, we found all the cells that ever fired a spike during the C segment of CR (set 1) or CL (set 2) replays. The size of intersection of the two sets was compared with that of either set to quantify overlap. Next, firing rates during the C segment of each joint replay were calculated for cells belonging to the union of the two sets. The firing rates of each cell in CR and CL replays were compared by using the Kolmogorov–Smirnov two-sample test. Last, the distribution of the absolute differences between cells' mean firing rates in CR and CL was compared with shuffles in which the types of the joint replays were randomized (group sizes were kept the same), also by using the Kolmogorov–Smirnov two-sample test.

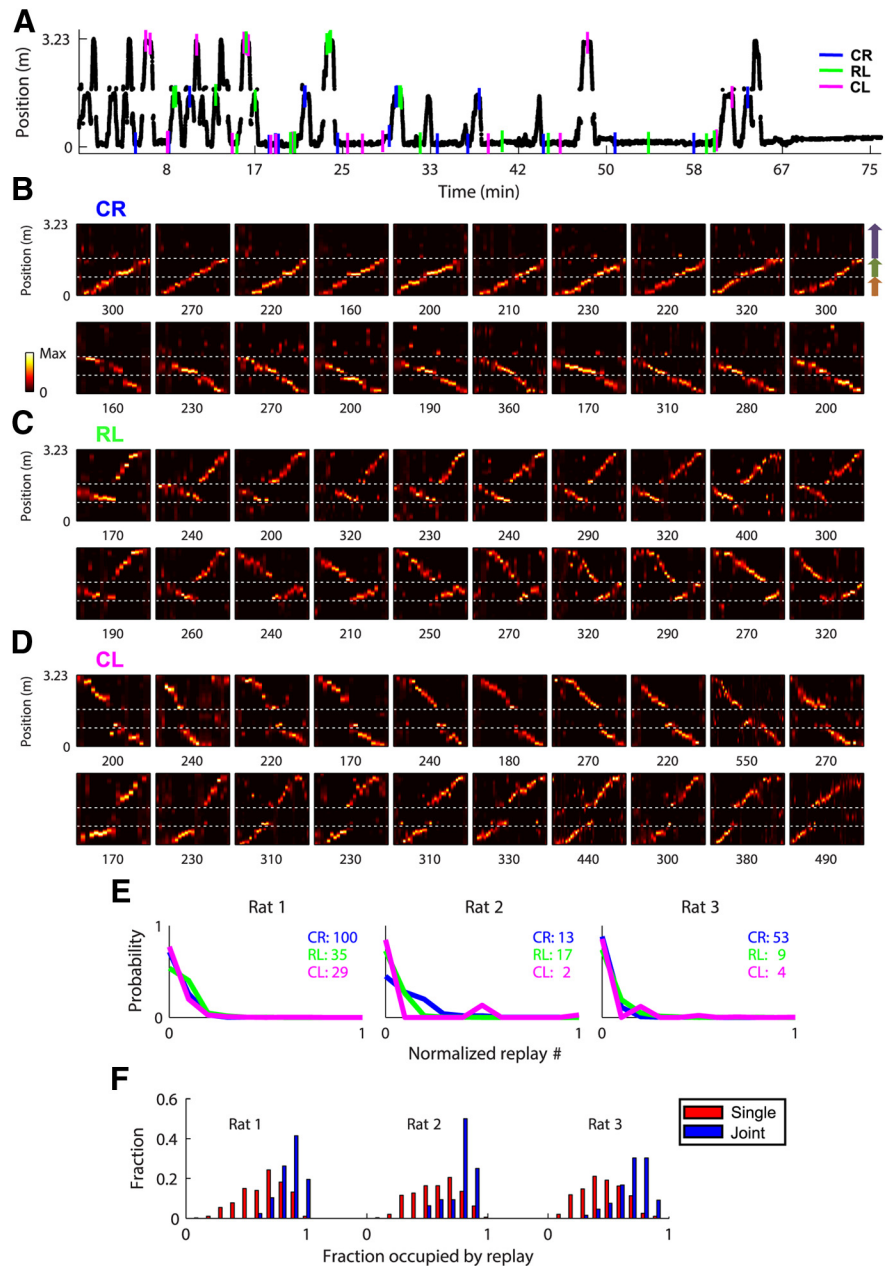
Due to the small numbers of joint replays of Rat 2 (Fig. 2E), we used only its CR and RL replays in the analyses in this section.

**Directionality of replay.** Directional place fields were calculated using spikes and position data from only inbound or only outbound laps for each arm. Joint probabilities estimated over both position and direction (Davidson et al., 2009; time bin = 10 ms) were computed for all identified replays. Two variables were calculated to quantify directionality by using joint probabilities within the corresponding trajectory-specific subregions.

The first variable, “directional,” was defined as follows:

$$\frac{1}{M} \sum_{i=1}^M \frac{\text{abs}\left(\sum_{j=1}^N w_{1_{ij}} - \sum_{j=1}^N w_{2_{ij}}\right)}{\sum_{j=1}^N w_{1_{ij}} + \sum_{j=1}^N w_{2_{ij}}},$$

where  $w_{1_{ij}}$  and  $w_{2_{ij}}$  are the joint probabilities of the inbound and outbound directions of pixel  $(i, j)$ , and measures on average how directional



**Figure 2.** Joint replays were detected from neuronal data in significant numbers. **A**, Linearized position as a function of time for Rat 1 during first exposure to the modified Y maze. Colored ticks mark when and where example replays in **B**, **C**, and **D** occurred. **B**, **C**, **D**, Examples of identified replay events representing CR, RL, and CL trajectories from a single rat (Rat 1), in which the position is decoded from neuronal spike trains in nonoverlapping 10 ms bins. Horizontal dashed lines mark arm boundaries. Replays in each row are ordered by the time of occurrence. The duration of each event (in milliseconds) is shown below each example. **E**, For each pair of joint arms, the probability of observing the number of identified joint replays by chance is expressed as the distribution of the numbers of replays representing the same joint arms sampled from 5000 shuffles as fractions of the number of joint replays actually observed (numbers in inset). Each type of replay from each animal was highly significant. **F**, Histograms of fractions of candidate event time windows occupied by single-arm replays and joint replays are plotted for each animal, normalized by the total numbers of replays. Single-arm replays ( $0.6 \pm 0.0$  across all rats) occupied significantly smaller fractions of spike density events than joint replays ( $0.8 \pm 0.0$  across all rats; Rat 1:  $t_{(755)} = -14.6$ ;  $p < 10^{-42}$ ; Rat 2:  $t_{(385)} = -5.8$ ;  $p < 10^{-7}$ ; Rat 3:  $t_{(267)} = -10.1$ ;  $p < 0.02$ ).

a replay is during each of its time bins. It has a range of 0–1, with 0 indicating totally not directional and 1 indicating totally directional. Replays with values  $\geq 0.3$  were considered directional, otherwise they were considered nondirectional. The threshold is fixed so that it is independent of the data (same below).

For a replay determined to be directional, the second variable “bias” defined by the following:

$$\frac{\sum_{i=1}^M \sum_{j=1}^N w1_{ij} - \sum_{i=1}^M \sum_{j=1}^N w2_{ij}}{\sum_{i=1}^M \sum_{j=1}^N w1_{ij} + \sum_{i=1}^M \sum_{j=1}^N w2_{ij}},$$

was then used to measure whether this replay as a whole was consistently biased toward representing either direction. It has a range of  $-1$  to  $1$ , with  $-1$  indicating pure representation of the outbound direction and  $1$  indicating pure representation of the inbound direction. Replays with values larger than  $0.3$  were considered to have stronger representations of the inbound direction over the outbound direction; those smaller than  $-0.3$  were considered to have stronger representations of the outbound direction over the inbound direction; and those between  $-0.3$  and  $0.3$  were considered to have mixed representations of both directions.

Finally, replays were determined to be forward if bias and motion were in the same direction, reverse if bias and motion were in opposite directions and mixed if bias was mixed. For joint replays, directionalities of their two segments were separately determined using joint probabilities within the two corresponding subregions.

**Ripple and multiunit activity analyses.** Ripple amplitude was calculated as in Davidson et al. (2009), with minor changes. We filtered the LFP signal from each selected channel in the ripple band ( $150$ – $250$  Hz), and Hilbert transformed the filtered signal to compute its envelope as the absolute value of the transformation. The mean envelope averaged across all selected tetrodes ( $15$ ,  $13$ , and  $13$  of  $19$  tetrodes for Rats 1, 2, and 3) was smoothed with a Gaussian window ( $SD = 8$  ms) to represent a continuous mean ripple amplitude. Individual ripples were also detected as local peaks in the ripple amplitude curve  $>2.5$  SDs above the mean, which were both calculated across all stopping periods. We defined the time point at which replay “passes” the choice point as the mean of inner boundaries of the two corresponding single-arm trajectory-specific subregions. The ripple amplitude trace associated with each joint replay was aligned to this time point to compare ripple amplitude between choice point representation and representations of the preceding and following arms across all joint replays. Multiunit spikes were defined as all recorded spikes whose largest amplitudes across tetrode channels were  $>100$   $\mu V$ . Multiunit spike density was smoothed across  $10$  ms bins (Gaussian filter;  $SD = 6$  ms).

## Results

Multitetrode recordings in the dorsal CA1 area of the hippocampus were conducted in three rats that were exploring the Y maze for the first time. Rats were allowed free exploration and were rewarded with chocolate milk at arm ends according to a spatial alternation rule (see Materials and Methods). Putative CA1 pyramidal single units were identified, and their place fields were computed. Only pyramidal cells with peak in-field firing rates exceeding  $1$  Hz ( $88$  from Rat 1,  $67$  from Rat 2,  $58$  from Rat 3) were used in the following analyses.

### Abundant joint replays spanning each two connecting arms were identified

To detect hippocampal replay, we first identified candidate events (mean duration =  $154.6$  ms) as transient increases in spike density across all cells, occurring during stopping periods restricted to the three reward areas. A Bayesian decoding algorithm (Davidson et al., 2009) was used to estimate posterior probabilities of position during candidate events. Based on posterior probabilities, candidate events were segmented in position and time into trajectory-specific subregions (Fig. 1C). Within each subregion, replays were defined as events with a high correlation between position and time, using a weighted correlation method with posterior probabilities as weights. We found large numbers of joint replays in all three animals, which extended across pairs of arms (CR, CL, and RL; Rat 1,  $N = 164$ ; Rat 2,  $N = 32$ ; Rat 3,  $N = 66$ ; see examples in Fig. 2B–D). These joint replays were

readily identified after just the first few running laps across arms and then consistently throughout the recording sessions (Fig. 2A, time points of example replays).

The significance of replay was evaluated by applying the replay identification method to shuffled data in which the individual spike train and place-field structure of each cell were preserved, but the spatial relationship between different cells was disrupted by circularly shifting the place field of each cell independently by a random number of position bins. For each rat, each type of joint replay was highly significant in number compared with shuffles [Monte Carlo method,  $p < 0.001$  (except for CL of Rat 2,  $p = 0.025$ , and CL of Rat 3,  $p = 0.002$ ); Fig. 2E].

We also observed large numbers of replays representing single arms (Rat 1,  $N = 593$ ; Rat 2,  $N = 355$ ; Rat 3,  $N = 203$ ). However, it is possible that many single-arm replays were partial joint replays where one of the arms was below the detection threshold. In support of this, the fraction of spike density events occupied by single-arm replays was significantly lower than for joint replays (Fig. 2F). As a result, we focused on joint replays in the following analyses, and also because the Y maze structure—the joining of three arms—could only be reflected in the joint replays.

### Multiple trajectories were replayed in the same stopping period

Further analyses of replay content revealed that across all stopping periods of the three animals  $90.8 \pm 2.4\%$  of the joint replays started from the current arm, confirming the initiation bias that has been reported (Foster and Wilson, 2006; Davidson et al., 2009). Previous reports of replay on a linear track demonstrated that each stopping period was associated with multiple replay events of the same trajectory (Foster and Wilson, 2006). This was interpreted as an efficient use of experience. Here we likewise found that individual stopping periods yielded multiple replay events ( $2.0 \pm 0.3$  joint replays, and  $8.9 \pm 1.5$  single-arm replays). However, in contrast to the linear track, we observed that on the Y maze individual stopping periods were associated with multiple replays depicting different trajectories. A large fraction of stopping periods that exhibited joint replay contained replay of multiple different trajectories ( $38.1\%$ ). Indeed, for stopping periods exhibiting joint replay, on average  $1.5 \pm 0.1$  types of joint-arm trajectory (from a range of  $1$ – $3$ ) were represented (Fig. 2A). Thus, stopping periods were associated with replays of more than one experience.

### Neuronal sequences were bifurcated

While the use of position estimation allowed the information content (i.e., trajectory) of replay sequences to be decoded, we further wished to identify the basis of this information content in the responses of individual neurons. In particular, we considered two models that might have accounted for the joint replays we observed. Joint replays could have been encoded using independent populations of neurons (Fig. 3A); thus, the underlying neuronal sequences would be linear, which would in essence be the same neuronal sequences encoding independent linear tracks. Alternatively, common arms of joint replays could have been encoded by the same cells (Fig. 3B); thus, the underlying neuronal sequences would be bifurcated, truly encoding the forked Y maze structure.

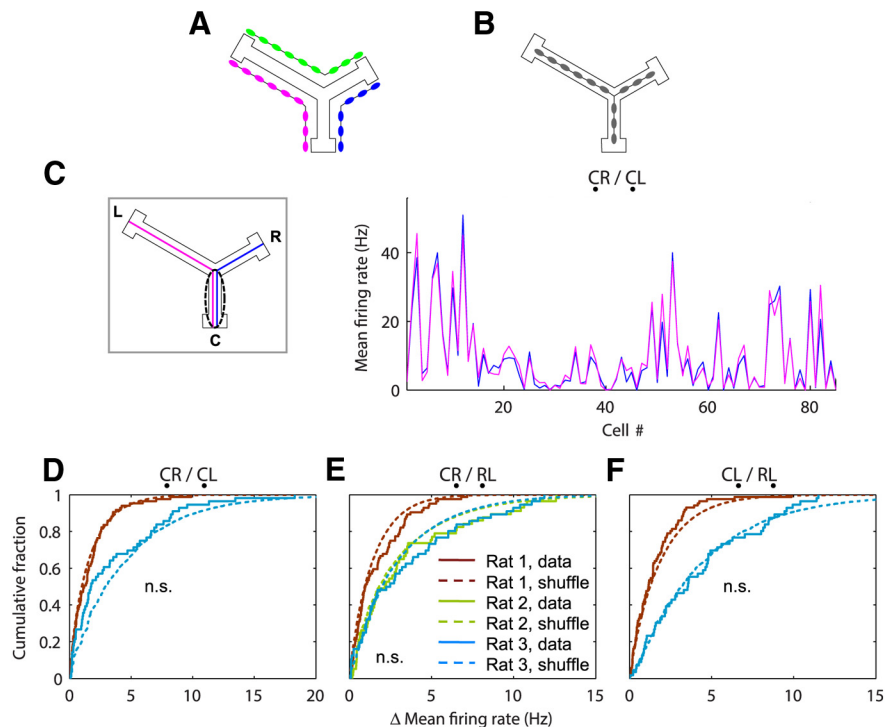
For each pair of joint replays proceeding from each arm, the cells that fired during the common segment were almost identical: on average,  $93.6 \pm 2.2\%$  of the cells that fired in one replay also fired in the other. Furthermore, these populations did not differ in firing rate between the two replays. For example, Figure

3C shows cells from Rat 1 that participated in firing during the C segments of CR or CL replays. Despite the wide range of firing rates among different cells, each individual cell had almost identical firing rates between CR and CL replays (the almost complete overlap between blue and magenta curves). Across all three rats, an average of  $98.0 \pm 0.5\%$  of the participating cells ( $94.9 \pm 1.7\%$  of all cells) had no significant difference between their firing rates during the common segments ( $p > 0.05$ , Kolmogorov–Smirnov two-sample test). Finally, the distributions of absolute differences in mean firing rate between the two replays were not significantly different from those of shuffles in which replay type was randomized ( $p > 0.10$  for all comparisons, Kolmogorov–Smirnov two-sample test; Fig. 3D–F). These results demonstrate that the same ensemble of cells fired during the common arm of overlapping joint replays with equal firing rates, as if they did not distinguish between the two types of replay, implying that joint replays reflected, at the neuronal level, the bifurcating spatial structure of the maze.

### Replays were highly directional

We next asked whether directionality (Fig. 4A) was encoded by replay and whether it might reflect the structure of the Y maze. Directional place fields were calculated and were used to compute joint posterior probabilities over both position and direction (Davidson et al., 2009), for the already identified replays. The directionality of a replay was quantified using two scores. “Directional” measured the extent to which the replay tended to be directional during each time bin regardless of which direction was preferred, on a scale of 0–1. “Bias” measured the extent to which a directional replay as a whole favored one direction over the other, on a scale of –1 (outbound) to 1 (inbound). We further compared bias with the direction of motion of replay to categorize replays as either forward (bias and motion in the same direction) or reverse (bias and motion in opposite directions). Across all three rats, the overwhelming majority of replay sequences (97.7%, including single-arm replays and the component segments of joint replays) had a directional score  $>0.3$ . This threshold corresponds to a divergence between the marginal probabilities in either direction such that the probability in one direction (0.65) was almost double that in the other (0.35). In fact, the mean directional score was  $0.62 \pm 0.00$ , which corresponds to a fourfold difference in marginal probabilities (0.81 to 0.19). Thus, replays were highly directional.

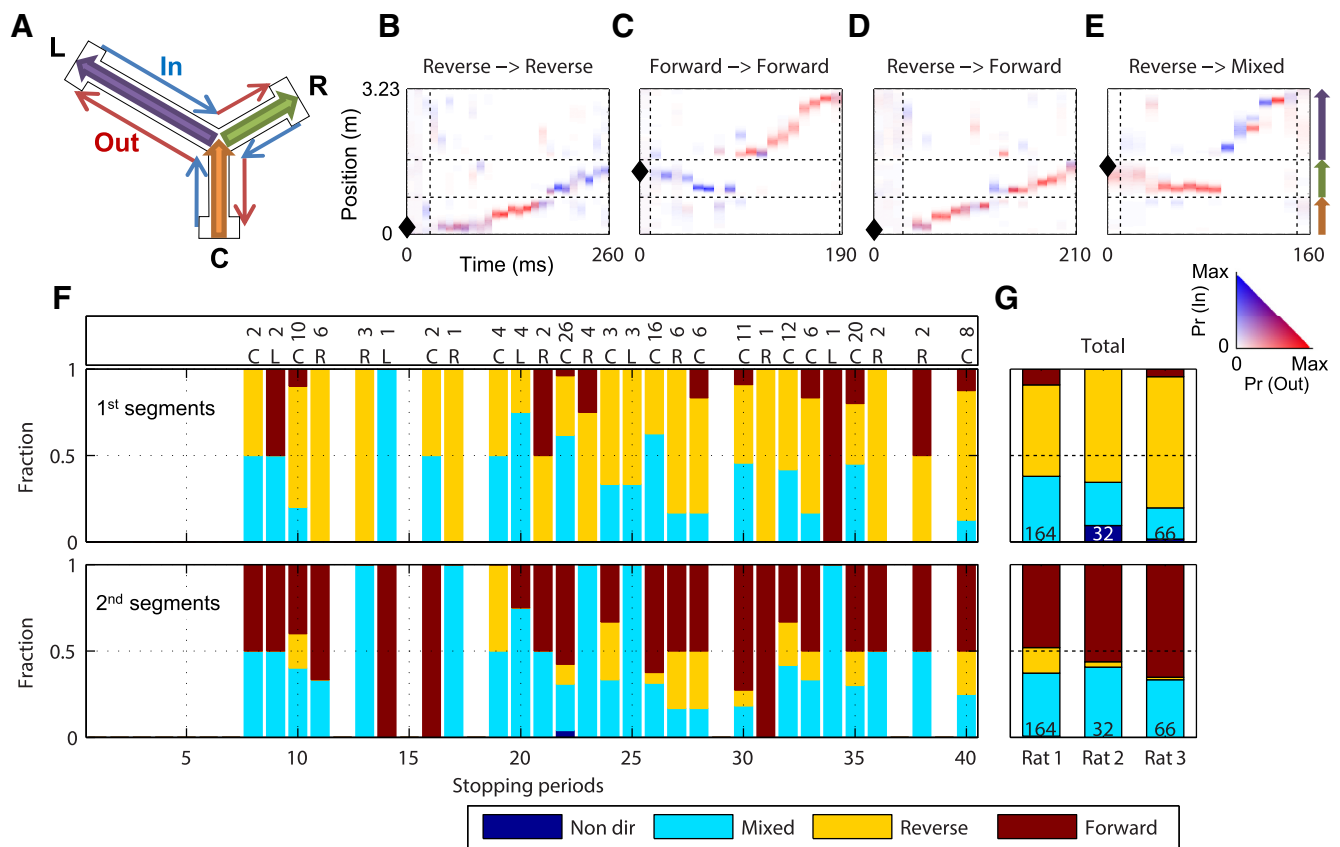
We then asked whether the components of joint replays exhibited the same bias. Interestingly, while joint replays exhibited various combinations of bias, only 8.0% were either consistently reverse (a reverse sequence of the outbound direction followed by a reverse sequence of the inbound direction, 4.6%; Fig. 4B) or consistently forward (a forward sequence of the inbound direction followed by a forward sequence of the outbound direction, 3.4%; Fig. 4C). By contrast, 31.3% of joint replays were composed



**Figure 3.** The same group of cells fired during common segments of joint replays. **A, B**, Two hypotheses of how joint replays might be encoded. **A**, Three independent populations of place cells with three separate sequences of place fields on the Y maze. Each pair of joint arms is encoded by a separate sequence. **B**, Joint replays were generated by a neuronal network that captures the spatial structure of the Y maze. The common segment of each pair of joint replays is generated by two different groups of cells in **A** and the same group of cells in **B**. **C**, Mean firing rates of cells from Rat 1 during the common segment C of joint CR (blue) and CL (magenta) replays. The included cells fired at least one spike during the C part of at least one CR or CL replay. Note that firing rates for the two replay types are highly similar across cells. **D, E, F**, Solid curves show cumulative distributions of absolute differences between the mean firing rates of cells during the common segments of paired joint replays (e.g., the blue and magenta curves in **C**). Dashed curves show distributions calculated from 5000 shuffles in which replay types (e.g., CR and CL) were randomized.

of segments with opposing biases (e.g., reverse in the outbound direction followed by forward in the outbound direction; Fig. 4D). Thus, a large fraction of joint replays switched directionality when “passing” the junctions of the two represented arms, implying the encoding of the choice point location in replay directionality.

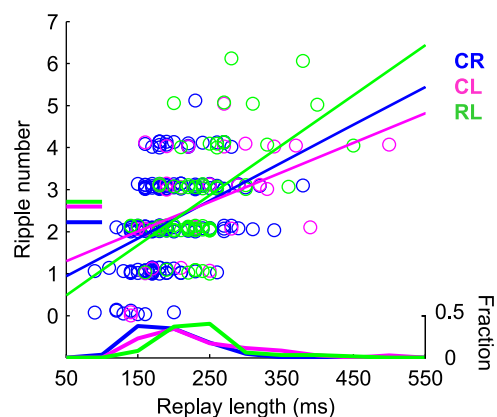
We then asked whether joint replays with opposing biases displayed random combinations or whether instead they obeyed an organizing principle. Strikingly, a far greater fraction of replays was reverse followed by forward (30.5%) rather than forward followed by reverse (0.8%). Comparing different directionalities for the joint replay segments separately, we found that consistently across stopping periods throughout a session, the first segment tended to be reverse (Fig. 4F, top) and the second segment tended to be forward (Fig. 4F, bottom). This consistency was maintained across replays initiated in different arms (Fig. 4F, overbar). This pattern of organization was consistent across the three animals. That is, for each of the three rats it was found that first segments were significantly more reverse than forward ( $p < 10^{-4}$  in each rat; Fig. 4G, top), whereas second segments were significantly more forward than reverse ( $p < 0.001$  in each rat; Fig. 4G, bottom). Given the bias for replay to start from the current location and proceed along multiple trajectories, these data suggest replays of the most immediate behavior were mostly reverse, while more diversified and distant replays were mostly forward.



**Figure 4.** Joint replay directionality. *A*, The junction of the three arms is the choice point. Running toward the choice point is “inbound,” and running away is “outbound.” *B–E*, Examples of joint replay with different combinations of directionalities from Rat 1. Horizontal dashed lines indicate arm boundaries. Black diamond shapes mark the location of the rat when each replay occurred. The color scale is set so that maximally saturated colors correspond to the largest position probability of each replay. *B*, A consistent reverse replay of CR. *C*, A consistent forward replay of RL. *D*, A CR replay with a reverse C segment followed by a forward R segment. *E*, An RL replay with a reverse R segment followed by a mixed L segment. *F*, Each bar shows, for all joint replays with at least one directional segment identified during each stopping period, the percentage of the first segments (top) or the second segments (bottom) with each directionality type (see legend). The joint replay number and location of each stopping period are shown in the overbar. Data are from Rat 1. *G*, Percentages calculated for all stopping periods combined, for the three rats separately. The total number of joint replays (with at least one directional segment) is shown in each bar plot for each rat. For each of the three rats, it was found that first segments were significantly more reverse than forward: Rat 1: reverse,  $0.58 \pm 0.06$ ; forward,  $0.13 \pm 0.05$ ;  $t_{(52)} = 6.0$ ;  $p < 10^{-6}$ ; Rat 2: reverse,  $0.67 \pm 0.11$ ; forward,  $0$ ;  $t_{(16)} = 6.1$ ;  $p < 10^{-4}$ ; Rat 3: reverse, fraction of reverse play =  $0.77 \pm 0.07$ ; forward,  $0.05 \pm 0.04$ ;  $t_{(52)} = 9.1$ ;  $p < 10^{-11}$ . For each of the three rats, it was also found that the second segments were significantly more forward than reverse: Rat 1: reverse,  $0.11 \pm 0.03$ ; forward,  $0.44 \pm 0.06$ ;  $t_{(52)} = -5.2$ ;  $p < 10^{-5}$ ; Rat 2: reverse,  $0.01 \pm 0.01$ ; forward,  $0.47 \pm 0.11$ ;  $t_{(16)} = -4.3$ ;  $p < 10^{-3}$ ; Rat 3: reverse,  $0.00 \pm 0.00$ ; forward,  $0.63 \pm 0.08$ ;  $t_{(52)} = -7.5$ ;  $p < 10^{-9}$ .

**Ripples were specifically associated with arms during joint replays**

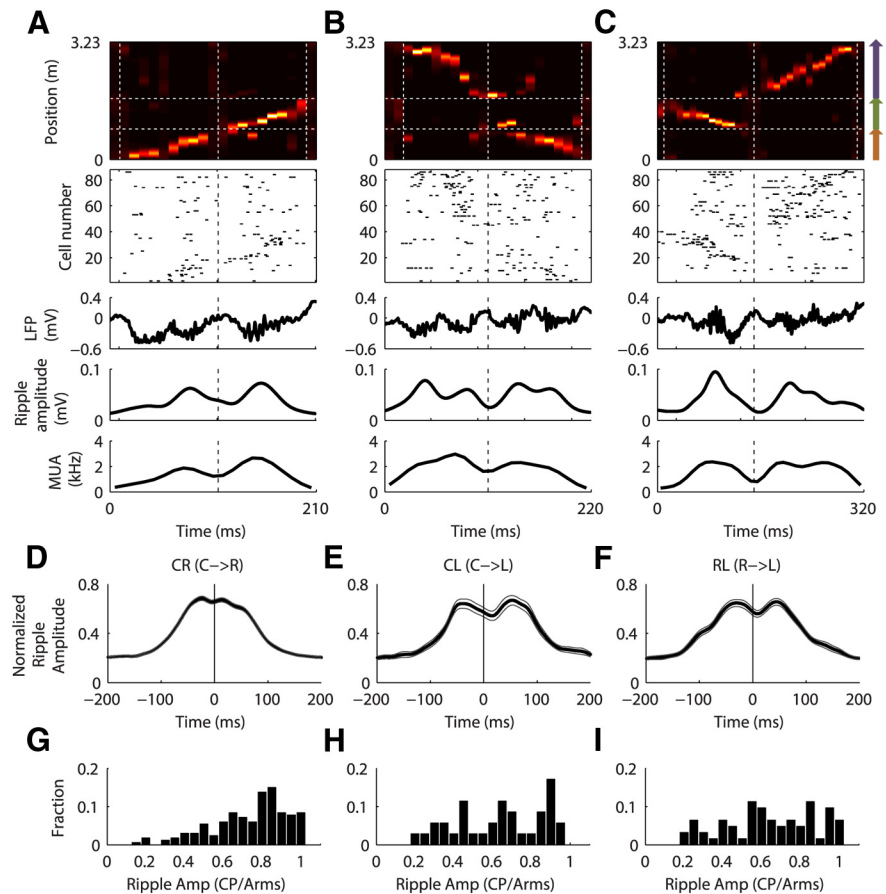
In a 10-m-long linear track, Davidson et al. (2009) observed extended replays covering several meters of the track, which were associated with multiple ripple events. Their findings suggested that replays in a large environment were not only extended in duration but furthermore were composed of discrete, shorter subevents. Considering the Y maze as a relatively large environment consisting of spatially distinct segments (the three arms), we wondered whether replay sequences would be associated with multiple ripple subevents, and whether these events would exhibit a correspondence with the maze structure. Indeed, a majority of joint replays (79.5%) were accompanied by more than one ripple event (median = two ripple events; Fig. 5). Moreover, joint replays representing the longer CL and RL trajectories were significantly longer and contained significantly more ripples than those representing the shorter CR trajectory (mean duration: CR,  $193.2 \pm 3.7$  ms; CL,  $234.9 \pm 13.7$  ms;  $t_{(200)} = -4.1$ ,  $p < 10^{-4}$ ; and RL,  $236.8 \pm 7.4$  ms;  $t_{(227)} = -5.7$ ,  $p < 10^{-7}$ ; mean ripple number: CR,  $2.2 \pm 0.1$ ; CL,  $2.6 \pm 0.2$ ;  $t_{(200)} = -1.8$ ,  $p = 0.068$ ; and RL,  $2.7 \pm 0.2$ ;  $t_{(227)} = -2.9$ ,  $p = 0.004$ , two-sided two-sample  $t$  tests),



**Figure 5.** Joint replays were associated with multiple ripples. The number of ripples detected during each joint replay (during the joint trajectory-specific subregion) was plotted against the duration (subregion length) of each replay, with each replay type plotted in a different color. For visualization purpose only, random noise was added to ripple numbers. Horizontal lines on left indicate mean ripple numbers. Histograms of replay durations are shown at bottom. Diagonal lines show linear regressions based on each replay type.

although not by 1.5 times, as were the physical joint-arm lengths. Each type of joint replay also exhibited a positive linear relationship between replay duration and associated ripple number (linear regression: CR:  $R^2 = 0.17$ , 9.0 ripples/s; CL:  $R^2 = 0.22$ , 7.0 ripples/s; RL:  $R^2 = 0.33$ , 11.9 ripples/s;  $p < 0.01$  for each replay type; compare with 9.9 ripples/s in the study by Davidson et al., 2009).

We then examined the relationship between ripples and the maze arms depicted in the associated replay sequences. Figure 6A–C shows example joint replays of each type along with underlying unit activities and accompanying LFP. It can be seen in the raw LFP recordings that one or two ripple events occurred during the replay of each individual arm, but ripples did not straddle the choice point. We calculated ripple amplitude as in Davidson et al. (2009) as a single trace averaging ripple activities across selected tetrodes, and observed similarly one to two peaks (corresponding to discrete ripple events) during individual arm representations, together with low ripple amplitude at choice point representations. We then aligned the ripple amplitude trace of each joint replay to the time point when choice point was represented, around which the mean traces across all CL or RL replays also showed “dips” in between two broad peaks (mean ratio of ripple amplitude at choice point to averaged amplitude between the highest peaks on either side of the choice point: CL,  $0.64 \pm 0.04$ ; RL,  $0.64 \pm 0.03$ ; Fig. 6E, F, H, I, offset of dips from 0 ms most likely is due to inaccuracies in “choice point passing time” calculations). Interestingly, the dip in the mean CR trace was very small (choice point-to-peak average ratio =  $0.74 \pm 0.02$ ; Fig. 6D, G), suggesting that the effect was most prominent for replays representing longer trajectories that extended for longer durations and contained more ripples. Furthermore, this effect was not simply due to weaker place-field representation at the choice point. By computing two-dimensional place fields, we compared the mean firing rate across all pixels located within choice-point area ( $\leq 10$  cm along linearized directions to track center), with the mean firing rate across all pixels within track arms ( $> 15$  cm along linearized directions to track center) across all cells. We found that place-field representation at the choice point was similar to, or marginally stronger than, the representation at arm portions of the Y maze (Rat 1: choice point,  $1.72 \pm 0.22$  Hz; track arms,  $1.33 \pm 0.14$  Hz;  $t_{(174)} = 1.5$ ;  $p = 0.14$ ; Rat 2: choice point,  $1.83 \pm 0.34$  Hz; track arms,  $1.12 \pm 0.13$  Hz;  $t_{(132)} = 1.9$ ;  $p = 0.055$ ; Rat 3: choice point,  $1.67 \pm 0.24$  Hz; track arms,  $1.34 \pm 0.14$  Hz;  $t_{(114)} = 1.2$ ;  $p = 0.23$ , unpaired two-sided two-sample  $t$  tests). Thus, in addition to demonstrating that extended joint replays were composed of multiple discrete subevents, we made the observation that these subevents clustered into representations of



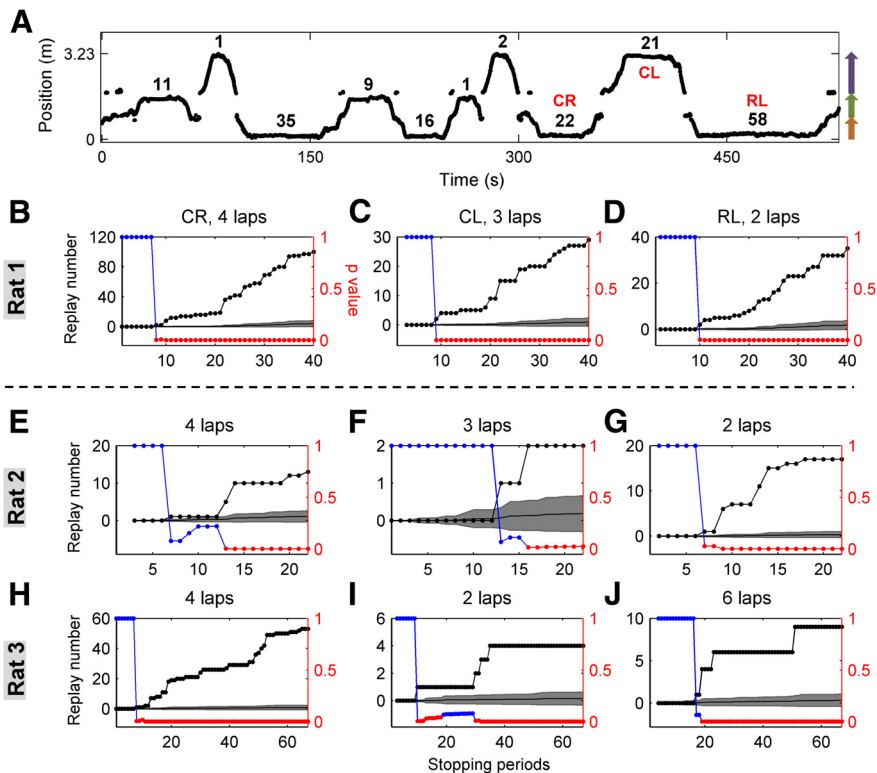
**Figure 6.** Ripples specifically co-occurred with arm representations. **A–C**, Examples from Rat 1. Top to bottom, Decoded joint replays of CR, CL, and RL; place-cell spikes during replay with cells ordered by locations of their peak firing rates on linearized track; raw LFP recording from one selected tetrode channel; ripple amplitude; multiunit spike density as a function of time. In all panels, the vertical dashed lines at the center indicate the time of choice-point representation during replay, which in **B** were moved to the left by two bins (20 ms) for illustration purposes solely. **D–F**, Ripple amplitude traces (e.g., those shown in **A–C**) of all joint replays, each normalized to its own maximum amplitude, were each aligned to the time when the choice point was represented during replay (0 ms in each panel, indicated by vertical black lines). They were also reoriented to the same joint-arm directions noted in captions (e.g., traces of R → C replays were all flipped around 0 ms). Mean  $\pm$  SEM values of the resulting traces across all three rats are shown for each replay type. **G–I**, Histograms of the ratios of ripple amplitude at 0 ms to mean of peak ripple amplitudes on either side of 0 ms.

entire single arms, rather than of random portions of the joint arms.

### Joint replays were detected after little experience

We further asked, given that joint replays with environment-specific structure occurred during the first exposure to the environment, how rapidly within the session can such replay be detected? To address this question, we determined the stopping periods by which joint replays reached significant numbers by counting the cumulative number of each type of joint replay for each stopping period, in original and shuffled data. We then counted the number of experiences of a replay before the first significant stopping period for that replay. For example, for Rat 1, all three types of joint replay rapidly outgrew shuffles in numbers (Fig. 7B–D) and reached significant numbers after two to four experiences of the corresponding joint arms (Fig. 7A). This pattern was replicated for every joint replay type, in each animal independently. Each animal acquired at least one of the joint replays after only two laps on the track (Fig. 7B–J). Across all animals, an average of  $3.3 \pm 0.4$  laps on corresponding joint arms were experienced before the first significant stopping periods.





**Figure 7.** Joint replays reached significant numbers rapidly. **A**, The first 10 stopping periods of Rat 1 (recorded position). The numbers of candidate events found during each stopping period are noted next to the stopping periods. Red letters indicate the stopping periods by which the noted types of replay first reached significant numbers. **B–D**, Each panel demonstrates how the number of observed replays outgrows those counted from shuffles (data from Rat 1). Dotted black lines, Cumulative numbers of replays detected from the original data. Gray shadings, Mean  $\pm$  SD of cumulative numbers of replays of the same type detected from 5000 sets of shuffled data. Colored dotted lines, Monte Carlo  $p$  values of the original cumulative numbers: blue,  $p \geq 0.05$  (not significant); red,  $p < 0.05$  (significant). The numbers of laps run on corresponding joint-arms before the first significant stopping periods are noted in titles. **E–J**, Results for Rats 2 and 3 are shown in **E–G** and **H–J**.

## Discussion

### Y maze structure was captured by individual neuron and population activities

In this study, we asked whether and how hippocampal place-cell sequences would reflect the topological structure of an environment of unpredictable shape, with the hypothesis that navigational useful sequences should capture this structure. We found that patterns of replay developed during the first exposure to a nonlinear environment that matched the sequential structure of the environment. This structure included unique elements such as bifurcated paths, and unequal lengths of the track arms. Most importantly, this nonlinear structure was captured at the level of the individual neurons. Spatially overlapping episodes (i.e., the common part of pairs of joint-arm traversals) were not replayed by independent populations of neurons, but by the same neurons with the same firing rates. Hence, the neuronal activities appear to be effectively “stitched together” in a manner reflecting the shape of the maze. Since the structure of the maze was unpredictable, these findings suggest that the neural network mechanisms responsible for generating these place-cell sequences involve learning.

The common coding scheme for spatially overlapping replays may have functional implications. Even though the animals traversed one joint-arm trajectory at a time, the separate running experiences were integrated in the brain so as to reflect the correct connectivity among the three track arms. This observed property of replay implies the construction of a map-like hippocampal

representation of the navigational environment as a whole, which importantly enables the “prediction of what leads to what” (O’Keefe and Nadel, 1978) and consequently flexible calculations of efficient routes between any two locations within the environment. This form of representation may support the generation of novel routes, such as short cuts, as well as appropriate generalization in the face of changes to the environment (e.g., effective detours to avoid novel obstacles). Such a scheme has clear advantages for navigation over the unintegrated representation of separate experiences, as would be the case if bifurcated paths were encoded by different populations of neurons, since this scheme would only support the selection of paths from the limited repertoire of encoded paths gained through direct experience. This observation made in the relatively simple Y-maze may imply a basic coding scheme for more generalized environments, whereby the hippocampus encodes the complete spatial structure, as opposed to separate running experiences.

Additionally, we found that joint replays were accompanied by multiple ripple events, confirming a recent finding that extended replays may be composed of discrete subevents (Davidson et al., 2009). Moreover, ripple occurrence was confined within the boundaries of joint replay segments that represented individual arms, which suggested that the precise times at which ripples were generated

during replay were not random. Instead, one or more subevents were closely clustered to represent linear components of the environment, while longer intervals between clusters accurately signaled the location of the arm intersection, again matching the spatial structure of the Y maze. Furthermore, the difference in arm lengths was also reflected by replay, in that physically longer trajectories (CL and RL) were represented by replays longer in duration, and correspondingly larger numbers of ripple events. The ratios of replay duration and ripple number between CL/RL replays and CR replays did not exactly match the ratios of track lengths (1.2 compared with 1.5), but this might have been due to systematic bias in the measurement of longer sequences when recording from limited numbers of hippocampal neurons.

### Implications of joint replay directionality

The directionality of replay sequences, and the existence of both forward and reverse sequences, have been reported previously for linear tracks and linear trajectories (Lee and Wilson, 2002; Foster and Wilson, 2006; Csicsvari et al., 2007; Diba and Buzsáki, 2007; Davidson et al., 2009). Here, we further show that replays of the linear pieces of a more complex environment are also overwhelmingly directional. Previous reports of replay on an extended track indicated that replays did not always represent a consistent directionality, but could flip between forward and reverse (Davidson et al., 2009). Here we also see direction flipping, which in fact predominated among joint replays, with the additional finding that replay tended to flip direction around the

choice point, supporting the notion that replay captured the unique structure of the environment. We further found that rather than the combination of forward and reverse components occurring at random, there is a distinct organizational pattern such that first segments of joint replays tended to be reverse, and second segments of joint replays tended to be forward. This finding combines in a fascinating way with two other findings in this task. First, we find that replay tends to start in the current arm. Second, we find that successive replays during single stopping periods can proceed along different trajectories. Integrating across these results, we can make the observation that joint replays tended to begin on the current arm and proceed in reverse order, before switching at the choice point to proceed along either of the two other arms in forward order. This organization suggests that reverse and forward replays may have different functions, with reverse replay representing a rewind of the immediate past, and forward replay representing the exploration of alternative futures, perhaps for the purposes of planning future behavioral trajectories. It is important to distinguish this classification of forward and reverse based on the directional tuning of place fields during bidirectional running (i.e., replays extending from point A to point B, or from point B to point A, can both be either reverse or forward), as used in the original report of reverse replay (Foster and Wilson, 2006), from an alternative classification that has been used, based on whether replay extends along the same unidirectional running path imposed by the task, or the opposite, never experienced running direction (i.e., replays extending from point A to point B are defined as forward, replays extending from point B to point A are defined as reverse; Gupta et al., 2010).

### Rapid occurrence of joint replays

Recurrent networks in the hippocampus have highly modifiable synapses, which are likely to undergo rapid synaptic plasticity during exploration (Buzsáki, 1989; Wilson and McNaughton, 1994; Martin et al., 2000; O'Neill et al., 2010), giving rise to rapidly reconfigured hippocampal circuits. Several computational models of the hippocampus have established the feasibility of using experience-dependent synaptic plasticity to acquire novel sequences (Jensen and Lisman, 1996; Levy, 1996; Leibold and Kempter, 2006; Molter et al., 2007; Koene and Hasselmo, 2008), in some cases after a single trial of behavioral experience. Consistent with these models, we found that joint replays were detected in significant numbers after very few trials of experience. For several reasons, these numbers of trials before significant replay was detected are likely to be an overestimate. First, stopping period durations were at the discretion of the animal, and longer stopping periods would have increased the probability of observing replay at earlier time points. Second, many single-arm replays may have been in fact joint replays for which one arm failed to pass the detection threshold (Fig. 2F) because of inherent experimental limitations on our ability to measure replay, given that the number of neurons recorded represents only a tiny fraction of the total network. Rapid learning is a prominent feature of hippocampally dependent learning (Morris, 2001), and so the rapid learning of replay sequences may play a fundamental role in hippocampal function.

Finally, while sequences can be encoded rapidly, they are also rapidly fit into a structure that captures the sequential spatial structure of the environment. Thus, while hippocampal replay may be formed out of the experience of individual episodes, its adaptive role may rather lie in the construction of predictive representations to guide future behavior (Wood et al., 1999;

Frank et al., 2000; Schacter and Addis, 2007; Pfeiffer and Foster, 2013). Determining the full relationship between hippocampal replay and hippocampally dependent learning, memory, and planning is a key future goal.

### References

- Buckner RL (2010) The role of the hippocampus in prediction and imagination. *Annu Rev Psychol* 61:27–48, C1–C8. [CrossRef Medline](#)
- Buzsáki G (1989) Two-stage model of memory trace formation: a role for “noisy” brain states. *Neuroscience* 31:551–570. [CrossRef Medline](#)
- Carr MF, Jadhav SP, Frank LM (2011) Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nat Neurosci* 14:147–153. [CrossRef Medline](#)
- Csicsvari J, O'Neill J, Allen K, 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. [CrossRef Medline](#)
- Davidson TJ, Kloosterman F, Wilson MA (2009) Hippocampal replay of extended experience. *Neuron* 63:497–507. [CrossRef Medline](#)
- Diba K, Buzsáki G (2007) Forward and reverse hippocampal place-cell sequences during ripples. *Nat Neurosci* 10:1241–1242. [CrossRef Medline](#)
- Ego-Stengel V, Wilson MA (2010) Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus* 20:1–10. [CrossRef Medline](#)
- Foster DJ, Knierim JJ (2012) Sequence learning and the role of the hippocampus in rodent navigation. *Curr Opin Neurobiol* 22:294–300. [CrossRef Medline](#)
- Foster DJ, Wilson MA (2006) Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* 440:680–683. [CrossRef Medline](#)
- Frank LM, Brown EN, Wilson M (2000) Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* 27:169–178. [CrossRef Medline](#)
- Gaffan D (1994) Scene-specific memory for objects: a model of episodic memory impairment in monkeys with fornix transection. *J Cogn Neurosci* 6:305–320. [CrossRef Medline](#)
- Gallistel CR (1990) *The organization of learning*. Cambridge, MA: MIT.
- Girardeau G, Benchenane K, Wiener SI, Buzsáki G, Zugaro MB (2009) Selective suppression of hippocampal ripples impairs spatial memory. *Nat Neurosci* 12:1222–1223. [CrossRef Medline](#)
- Gupta AS, van der Meer MA, Touretzky DS, Redish AD (2010) Hippocampal replay is not a simple function of experience. *Neuron* 65:695–705. [CrossRef Medline](#)
- Jadhav SP, Kemere C, German PW, Frank LM (2012) Awake hippocampal sharp-wave ripples support spatial memory. *Science* 336:1454–1458. [CrossRef Medline](#)
- Jensen O, Lisman JE (1996) Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. *Learn Mem* 3:279–287. [CrossRef Medline](#)
- Ji D, Wilson MA (2007) Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat Neurosci* 10:100–107. [CrossRef Medline](#)
- Karlsson MP, Frank LM (2009) Awake replay of remote experiences in the hippocampus. *Nat Neurosci* 12:913–918. [CrossRef Medline](#)
- Koene RA, Hasselmo ME (2008) Reversed and forward buffering of behavioral spike sequences enables retrospective and prospective retrieval in hippocampal regions CA3 and CA1. *Neural Netw* 21:276–288. [CrossRef Medline](#)
- Lee AK, Wilson MA (2002) Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* 36:1183–1194. [CrossRef Medline](#)
- Leibold C, Kempter R (2006) Memory capacity for sequences in a recurrent network with biological constraints. *Neural Comput* 18:904–941. [CrossRef Medline](#)
- Levy WB (1996) A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus* 6:579–590. [CrossRef Medline](#)
- Louie K, Wilson MA (2001) Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* 29:145–156. [CrossRef Medline](#)
- Martin SJ, Grimwood PD, Morris RG (2000) Synaptic plasticity and memory: an evaluation of the hypothesis. *Annu Rev Neurosci* 23:649–711. [CrossRef Medline](#)

- Molter C, Sato N, Yamaguchi Y (2007) Reactivation of behavioral activity during sharp waves: a computational model for two stage hippocampal dynamics. *Hippocampus* 17:201–209. [CrossRef Medline](#)
- Montague PR, Dayan P, Sejnowski TJ (1996) A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J Neurosci* 16:1936–1947. [Medline](#)
- Morris RG (2001) Episodic-like memory in animals: psychological criteria, neural mechanisms and the value of episodic-like tasks to investigate animal models of neurodegenerative disease. *Philos Trans R Soc Lond B Biol Sci* 356:1453–1465. [CrossRef Medline](#)
- Nakazawa K, Quirk MC, Chitwood RA, Watanabe M, Yeckel MF, Sun LD, Kato A, Carr CA, Johnston D, Wilson MA, Tonegawa S (2002) Requirement for hippocampal CA3 NMDA receptors in associative memory recall. *Science* 297:211–218. [CrossRef Medline](#)
- O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* 34:171–175. [CrossRef Medline](#)
- O'Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. London: Clarendon.
- Olton DS, Samuelson RJ (1976) Remembrance of places past: spatial memory in rats. *J Exp Psychol Anim Behav Process* 2:97–116.
- O'Neill J, Pleydell-Bouverie B, Dupret D, Csicsvari J (2010) Play it again: reactivation of waking experience and memory. *Trends Neurosci* 33:220–229. [CrossRef Medline](#)
- Pfeiffer BE, Foster DJ (2013) Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* 497:74–79. [CrossRef Medline](#)
- Schacter DL, Addis DR (2007) The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Philos Trans R Soc Lond B Biol Sci* 362:773–786. [CrossRef Medline](#)
- Scoville WB, Milner B (1957) Loss of recent memory after bilateral hippocampal lesions. *J Neurol Neurosurg Psychiatry* 20:11–21. [CrossRef Medline](#)
- Singer AC, Frank LM (2009) Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron* 64:910–921. [CrossRef Medline](#)
- Steele RJ, Morris RG (1999) Delay-dependent impairment of a matching-to-place task with chronic and intrahippocampal infusion of the NMDA-antagonist D-AP5. *Hippocampus* 9:118–136. [CrossRef Medline](#)
- Tolman EC (1948) Cognitive maps in rats and men. *Psychol Rev* 55:189–208. [CrossRef Medline](#)
- Wilson MA, McNaughton BL (1994) Reactivation of hippocampal ensemble memories during sleep. *Science* 265:676–679. [CrossRef Medline](#)
- Wood ER, Dudchenko PA, Eichenbaum H (1999) The global record of memory in hippocampal neuronal activity. *Nature* 397:613–616. [CrossRef Medline](#)