

# Hippocampal place-cell sequences depict future paths to remembered goals

Brad E. Pfeiffer<sup>1</sup> & David J. Foster<sup>1</sup>

**Effective navigation requires planning extended routes to remembered goal locations. Hippocampal place cells have been proposed to have a role in navigational planning, but direct evidence has been lacking. Here we show that before goal-directed navigation in an open arena, the rat hippocampus generates brief sequences encoding spatial trajectories strongly biased to progress from the subject's current location to a known goal location. These sequences predict immediate future behaviour, even in cases in which the specific combination of start and goal locations is novel. These results indicate that hippocampal sequence events characterized previously in linearly constrained environments as 'replay' are also capable of supporting a goal-directed, trajectory-finding mechanism, which identifies important places and relevant behavioural paths, at specific times when memory retrieval is required, and in a manner that could be used to control subsequent navigational behaviour.**

A fundamental purpose of memory lies in using previous experience to inform current choices, directing behaviour towards reward and away from negative consequences based on knowledge of prior outcomes in similar situations. Goal-directed spatial navigation—planning extended routes to remembered locations—requires both memory of the goal location and knowledge of the intervening terrain to determine an efficient and safe path. The hippocampus has long been known to have a critical role in spatial memory<sup>1,2</sup> and memory for events<sup>3,4</sup>, and it has been proposed that the hippocampus may have a fundamental role in calculating routes to goals, especially under conditions demanding behavioural flexibility<sup>1,5–8</sup>. This proposal stems largely from the discovery that excitatory neurons of the hippocampus show spatially localized place responses during exploration<sup>1</sup>. However, it has been a challenge to understand how individual place responses tied to the current location might be informative about other locations that the animal cares about, such as the remembered goal<sup>9</sup>, or the set of locations defining a route<sup>10,11</sup>.

Techniques to record simultaneously from multiple hippocampal place cells<sup>12</sup> have been used to show that place cells systematically represent positions other than the current location. The early discovery of phase precession of place-cell spikes relative to theta frequency oscillations in the local field potential<sup>13</sup> led to the hypothesis that place cells fire in sequences within a theta cycle, and thus represent places behind or ahead of the animal<sup>14–16</sup>. Theta sequences have since been demonstrated experimentally across place-cell populations<sup>17</sup>. Also during theta, place-cell activity seems to 'sweep' ahead of an animal located at a choice point<sup>18</sup>, leading to the hypothesis that such activity could support the evaluation of alternatives during decision making<sup>19</sup>. A separate group of phenomena termed 'replay' has been found during sleep<sup>20,21</sup> and non-exploratory awake periods<sup>22</sup>, and is associated with sharp-wave-ripple (SWR) events in the hippocampal local field potential (with the sole exception of replay during rapid eye movement sleep<sup>20</sup>). In replay, simultaneously recorded populations of place cells show reactivation of temporal sequences reflecting prior behavioural trajectories up to 10-m long<sup>23</sup>. Although these forms of non-local activity are now well established<sup>17,23–26</sup>, it has proven difficult to establish a predictive relationship between non-local place-cell activity and behaviour<sup>18,26</sup>, because of the twofold technical problem of ensuring

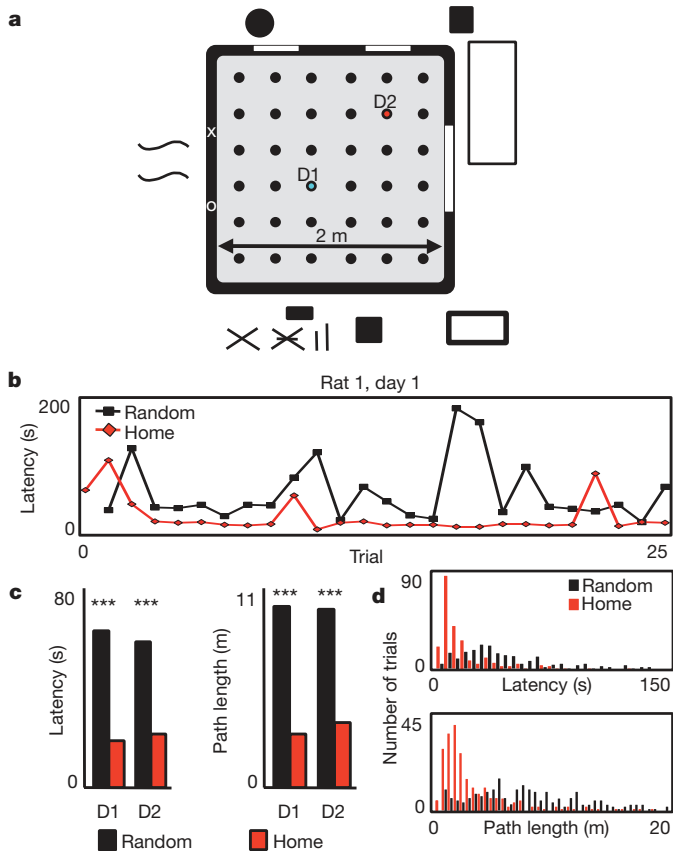
adequate behavioural sampling of the environment while recording from sufficient numbers of place cells. Thus it remains unknown whether non-local place-cell activity can specify remembered goals, or define specific routes that the animal will take.

## Depiction of two-dimensional trajectories

We recorded from hippocampal neurons while rats performed a spatial memory task, using the statistical power of an open-field design in which the goal was one of 36 clearly separated locations within a 2 m × 2 m arena (Fig. 1a). We addressed the sampling problem by combining random foraging and goal-directed behaviour, and by implanting miniaturized lightweight microdrives supporting 40 independently adjustable tetrodes, with 20 tetrodes targeted to each dorsal hippocampal area CA1 (Supplementary Fig. 1), to record simultaneously from up to 250 hippocampal neurons with well-defined place fields. Our task, incorporating elements from previous task designs<sup>9,27–29</sup>, was composed of trials each consisting of two phases: in phase one, the rat was required to forage to obtain reward (liquid chocolate) in an unknown location (Random). In phase two, the rat could obtain reward in a predictable reward location (Home). The transition to the next phase or trial was automatic upon consumption of the reward, and was not signalled to the animal. The task incorporated several features. First, because the shortest routes in phase one and two were matched, it was determined that animals could remember Home, but could not detect Random locations, because latencies and path lengths were significantly shorter for Home-bound trajectories (Fig. 1b–d). Second, the Home location was moved to a new location each day. Thus, animals were required to learn a new goal location, demanding a flexible behavioural response that was more likely to engage the hippocampus than a fixed reference-memory response<sup>27,30,31</sup>. Third, for the first 19 trials of each day, the Random locations were non-repeating. Hence during this period, every Home-bound trajectory was always a novel combination of current location and goal location. Thus, our task probed both memory for the goal location and flexible planning of a novel route to get there.

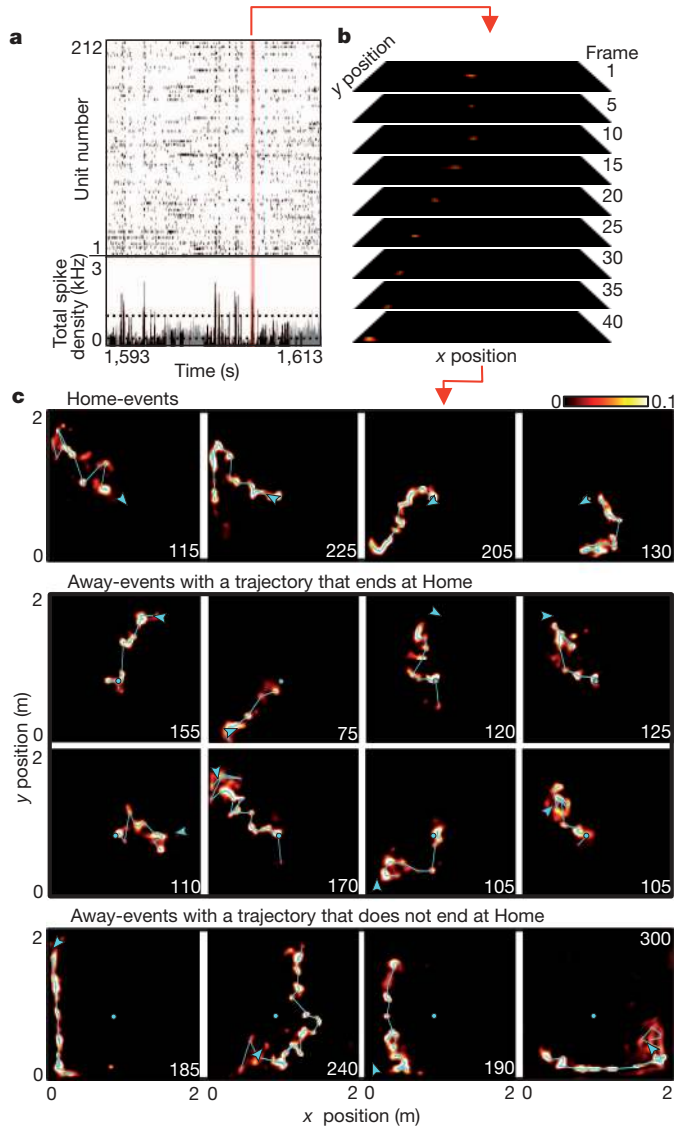
We implanted four well-trained rat subjects with the 40-tetrode microdrive for electrophysiological recording. Large numbers of

<sup>1</sup>Solomon H Snyder Department of Neuroscience, Johns Hopkins University School Of Medicine, Baltimore, Maryland 21205, USA.



**Figure 1 | Behaviour in the open-field spatial memory task.** **a**, Schema of arena and room, reward wells (circles), and Home location for days 1 and 2 (D1, cyan; D2, red). **b**, Per-trial latency to reach Home or Random well location across all rats for D1 and D2. **c**, Mean latency and path length to reach Home or Random well location across all rats for D1 and D2. **d**, Histogram of latencies (5-s bins) and path lengths (50-cm bins) for all trials (shown to 150 s and 20 m, respectively) **P**-values (Kolmogorov–Smirnov test): latency  $2.6 \times 10^{-2}$ ; path  $9.1 \times 10^{-4}$ .

well-isolated units (Supplementary Fig. 2)<sup>32</sup> were recorded simultaneously during behavioural sessions on two consecutive days (212 and 250 units active during exploration from rat 1 on experimental days 1 and 2, respectively; 166 and 193 units from rat 2 on days 1 and 2; 133 and 106 units from rat 3 on days 1 and 2; 103 and 175 units from rat 4 on days 1 and 2). The recorded units demonstrated position-specific firing patterns ('place fields') that were distributed throughout the environment (Supplementary Figs 3–5), and a memory-less, uniform prior Bayesian decoding algorithm<sup>23</sup> allowed us to estimate the spatial location of the rat accurately from the recorded spike trains throughout the experiment (Supplementary Fig. 6 and Supplementary Video 1). We identified candidate events as brief increases in population spiking activity during periods of immobility while the rat performed the task (Fig. 2a) and applied the decoding algorithm to the population spike trains (Fig. 2b). During many candidate events, decoded position revealed temporally compressed, two-dimensional trajectories across the environment (Fig. 2c and Supplementary Video 2). We applied length, duration and smoothness criteria to the decoded positions of candidate events to define 'trajectory events' (see Methods). We found between 144 and 373 trajectory events per session (between 25.3% and 43.9% of candidate events) with a mean duration of 103.6 ms, and path lengths that ranged from 40.0 cm to 199.1 cm (Supplementary Fig. 7 and Supplementary Table 1). We tested the probability that trajectory events could have occurred by chance, using two separate Monte-Carlo shuffle methods which varied either cell identity or place field position



**Figure 2 | Trajectory events.** **a**, Raster plot (top) and spike density (bottom) of simultaneous unit activity for R1,D1 for representative epoch. Periods of immobility denoted in black. Dashed lines represent candidate event detection threshold. **b**, Position posterior probabilities in selected frames for the candidate event in **a**. **c**, 16 representative events (of 274) for R1,D1, decoded and summed across time. Values indicated by colour bar. Event duration (in ms) in right corner. Cyan circle, Home well. Cyan line, peak probability for each timeframe. Cyan arrowhead, position and head direction of rat at time of event. Videos of each event available in online Supplementary Video 2.

(see Methods). Zero (out of 2,028) trajectory events had a *P*-value greater than 0.02 under either method, indicating that all trajectory events were statistically significant events. Spectrogram analysis of trajectory events strongly matched SWR events identified within the same experimental sessions (Supplementary Fig. 8a). In addition, an overwhelming majority of trajectory events were coincident with SWR events (Supplementary Fig. 8b). Theta power, which is high during exploration, was significantly decreased immediately before and after trajectory events (Supplementary Fig. 8c). Collectively, these data indicate that trajectory events are functionally similar to the SWR-associated events previously reported on linear tracks as 'replay'<sup>21–26</sup>.

### Trajectory events over-represent the goal

To examine whether non-local spatial information present in trajectory events contributes to or is affected by acquisition or expression of

a spatial memory (the novel Home location), we divided the observed trajectory events into those that were initiated while the rat was at the Home location ('home-events') and those that were initiated while the rat was elsewhere ('away-events'). There was no difference in the rate of occurrence of sharp-wave/ripple events or of trajectory events between Home and Random locations (Supplementary Figs 9 and 10). As expected, home-events showed strong representation of the Home location (Fig. 3a, c and Fig. 2c, top row), probably owing to initiation bias, a tendency for hippocampal events to reflect a path that begins at the rat's current location<sup>22–24</sup> (but see refs 25, 26). Strikingly, we observed that away-events also showed an increased representation of the Home location (Fig. 3b, d and Supplementary Fig. 11), a finding that cannot be explained through initiation bias. Consistent with this observation, many away-events depicted a trajectory that ended at Home (Fig. 2c, middle rows; Supplementary Videos 3–7). Quantification confirmed that the Home location was significantly over-represented in away-events relative to other locations on the

open field (Fig. 3e, left; Supplementary Fig. 12) and that away-events were more likely to end their trajectories at the Home location than any other region of the arena (Fig. 3e, centre).

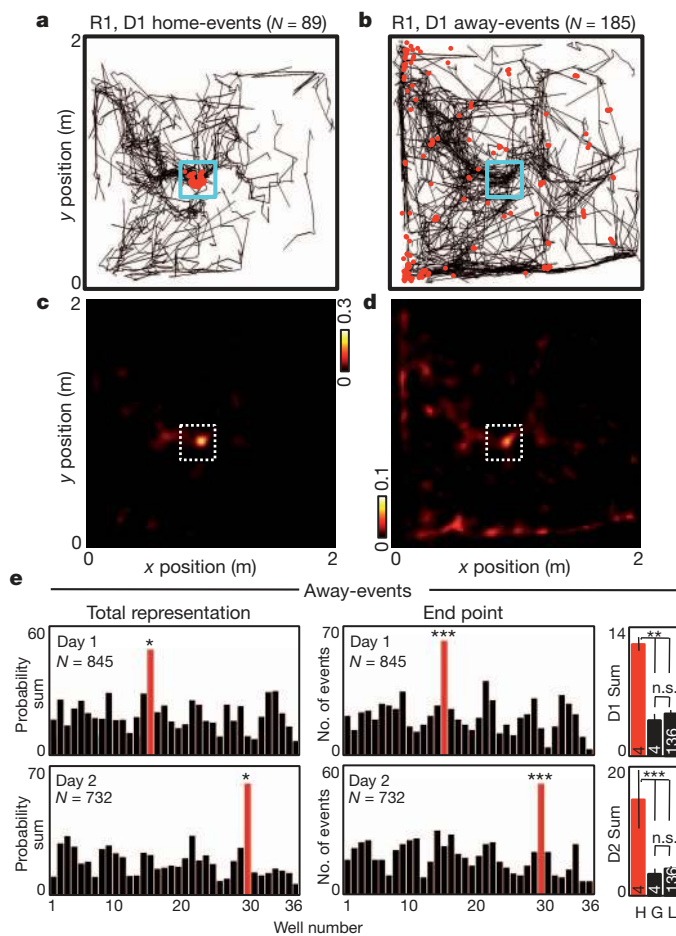
Importantly, the region of increased representation changed accordingly when the location of the Home well was moved on experimental day 2. The heightened representation of Home in away-events was present even when the analysis was restricted to the first 19 trials, when the specific Random–Home combinations were novel (Supplementary Fig. 13). The increased representation of Home in away-events was not a simple function of increased familiarity with or time spent at the Home location, as other regions of the arena with greater occupancy times did not show strong representations in trajectory events (Fig. 3e, right). The overexpression of the Home location in away-events could not be accounted for by either occupancy time or the spatial distribution of place fields (Supplementary Figs 14 and 15). Further, when we restricted our analysis to vectorized trajectories rather than entire posterior probabilities, the Home location remained over-represented in away-events (Supplementary Fig. 16). Thus, trajectory events in the hippocampus over-represent a known goal location in a manner which cannot be explained solely by occupancy time or place-field representation.

### No over-representation of non-goals

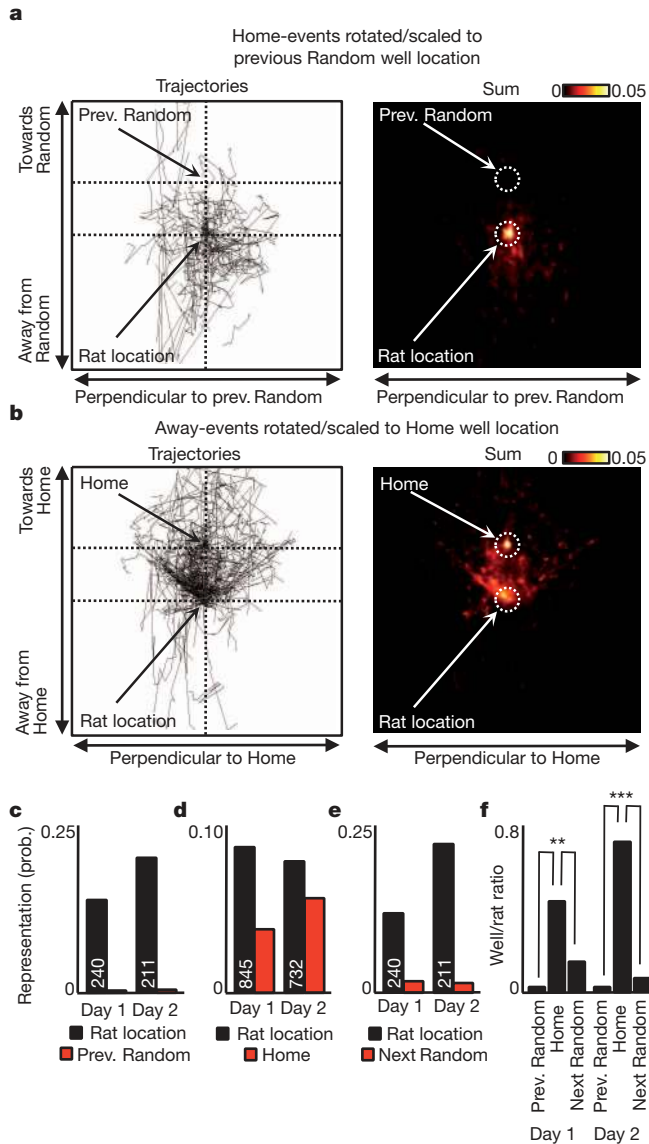
We proposed that the over-representation of locations in trajectory events was selective for behaviourally relevant locations. The task was designed so that the previous Random well was never a correct behavioural goal, and so we proposed in particular that the previous Random well would not be over-represented in trajectory events. To equalize comparison between away-events and home-events, we rotated and scaled all home-events such that the distance and direction from the rat's physical location at the time of each event to the previously active Random location was the same across all home-events (Fig. 4a). Similarly, we rotated and scaled all away-events according to the direction and distance to the Home location (Fig. 4b), and as a control we rotated and scaled all home-events according to the direction and distance to the immediately future (but unknown and not yet baited) Random well location. All rotated/scaled trajectory events showed a strong representation of the rat's physical location (Fig. 4c–e) due to initiation bias. However, whereas the rotated/scaled away-events showed a strong representation of the Home location (Fig. 4d), rotated home-events showed little representation of the previously active (Fig. 4c) or immediately-to-be active (Fig. 4e) Random locations. Indeed, we observed a significant decrease in the representation of the previous Random location in home-events compared to the representation of the Home location in away-events (Fig. 4f). These data show that hippocampal trajectory events reflect the demands of the task by selectively over-representing the immediately relevant Home location and not the irrelevant previous Random location.

### Trajectory events reflect future behavioural path

The initiation and termination bias that we observed suggested that away events depict the future trajectory to Home, indicative of a planning mechanism to guide behaviour. To test this hypothesis, we quantified the correspondence between trajectory events and the behavioural path in the immediate future, or immediate past (Fig. 5a, b and Supplementary Fig. 17). We calculated the angular displacement between trajectory and path at progressively increasing radii from the current location (Fig. 5a, b). Away-events were strongly concentrated around zero angular displacement assessed against the future path, and more broadly distributed with respect to the past path (Fig. 5c), and this difference was verified in terms of the mean absolute angular displacement for each event (Fig. 5d). Home events showed a weaker representation of future path, and an apparent anti-correlation with past path, which might have reflected the fact that the path back to the previous Random well was never correct (Fig. 5e, f). Away-events were significantly closer to the rat's future path than were home-events

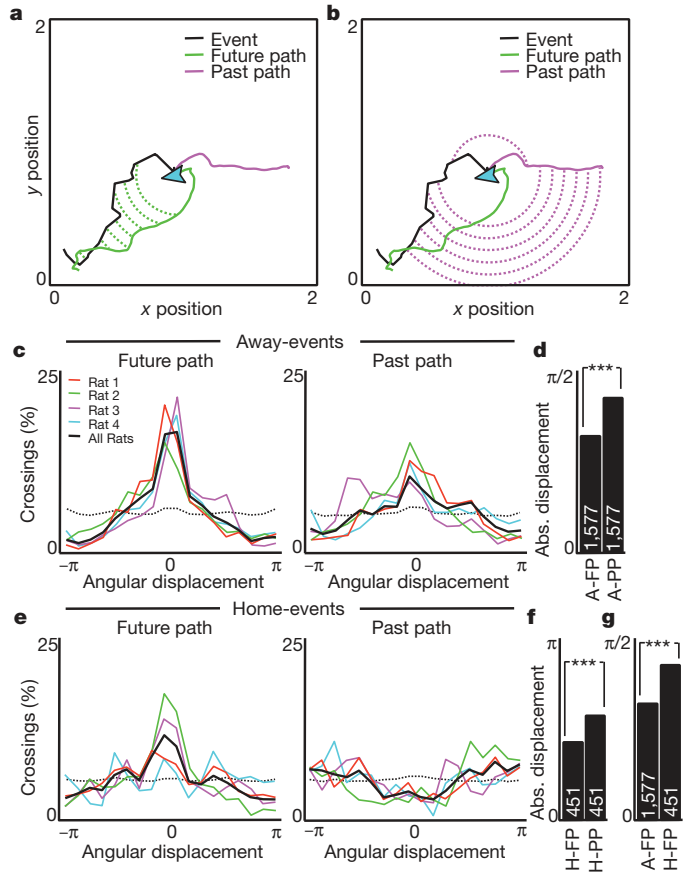


**Figure 3 | Remote representation of goal location.** a–d, Vectorized trajectories (a, b) and average posterior probability sum (c, d) of all confirmed home-events (left) and away-events (right) for R1, D1. Red dots in a, b, rat location at time of event. Dashed box in c, d, Home location. e, Left, posterior probability sum for all away-events across all rats. Home (red) is a statistical outlier. *P*-value (Grubbs' test for outliers): D1  $2.3 \times 10^{-2}$  (Lilliefors test, *P*-value 0.15); D2  $1.1 \times 10^{-2}$  (Lilliefors *P*-value 0.32). Centre, number of away-events across all rats in which the final frame peak posterior probability was at each well. Home (red) is a statistical outlier. *P*-values (Grubbs' test for outliers): D1  $6.9 \times 10^{-4}$  (Lilliefors test, *P*-value 0.29); D2  $6.0 \times 10^{-4}$  (Lilliefors *P*-value 0.42). Right, as left, but mean  $\pm$  s.e.m. for Home (H), all wells with greater in-session total occupancy than Home (G), and all wells with less occupancy than Home (L). *P*-values (ANOVA, Tukey–Kramer post-hoc multiple comparison): D1 H vs G  $2.9 \times 10^{-3}$ , H vs L  $8.5 \times 10^{-5}$ , G vs L 0.91; D2 H vs G  $7.4 \times 10^{-8}$ , H vs L  $1 \times 10^{-10}$ , G vs L 0.82.



**Figure 4 | Representation of relevant vs. irrelevant locations.** **a**, Vectorized trajectories (left) and average posterior probability sum (right) of all home-events for R1,D1, centred by the rat's physical location at time of event and rotated and scaled according to direction and distance to the previously rewarded Random location. White circles, quantified regions. Prev., previous. **b**, As **a**, for Home. **c–e**, Across all rats, mean representation of quantified regions as in **a**, **b**. Event number displayed on bar. **f**, Normalized ratio of well/rat representation for **c–e**. *P*-values (Wilcoxon rank-sum test): D1 Home vs prev. Random  $4.4 \times 10^{-16}$ , Home vs next Random  $9.9 \times 10^{-3}$ ; D2 Home vs prev. Random  $3.1 \times 10^{-20}$ , Home vs next Random  $1.3 \times 10^{-13}$ .

(Fig. 5g), consistent with the goal-directed nature of Random-to-Home navigation. We conducted two further analyses of path correspondence, one based on the orientation of the depicted trajectory to a location occupied 10 s in the future or the past (Supplementary Fig. 18), and one based on the spatial overlap between smoothed versions of the trajectory and future or past path (Supplementary Fig. 19), with matching results. Rats showed no bias to face the direction of their immediately future path or the Home well location during away-events (Supplementary Fig. 20a, b). Furthermore, away-events were more spatially correlated with the rat's future path than with his current heading (Supplementary Fig. 20d–g). Thus, the strong reflection of the rat's future path in away-events could not be trivially explained as a representation of paths 'in front' of the rat, but rather suggested a more precise path-finding mechanism.

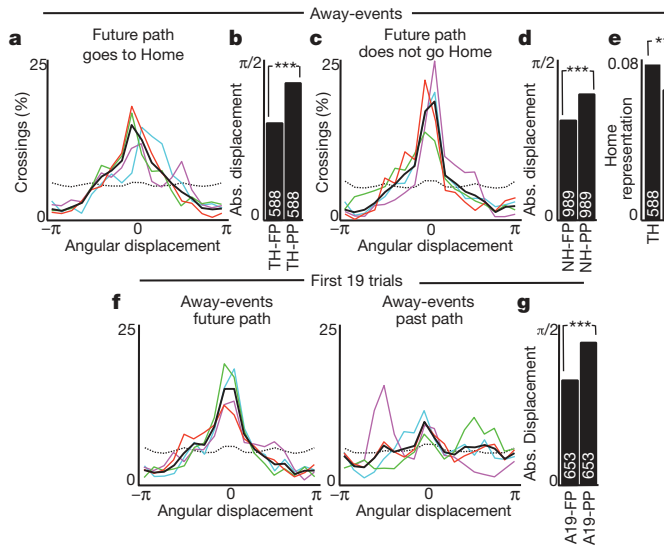


**Figure 5 | Correspondence to past or future path.** **a**, **b**, Representative event from R1,D1, demonstrating trajectory event vector (black), immediate future (green) and past (magenta) path (up to 10 s or 50 cm, whichever is greater), and angular displacement along the minor arc between event and future (**a**) or past (**b**) path at each crossing. **c**, Per cent of crossings across all events as a function of angular displacement for all away-events compared to future (left) or past (right) path. Dashed line indicates chance based upon 2,000 shuffled events. **d**, Mean absolute angular displacement for away-events compared to future (A-FP) or past (A-PP) path. Abs., absolute. **e**, **f**, As **c**, **d**, for home-events. **g**, Mean absolute angular displacement for future path for all away-events (A-FP) or home-events (H-FP). *P*-values (Wilcoxon rank-sum test):  $8.60 \times 10^{-31}$  (**d**);  $3.54 \times 10^{-17}$  (**f**);  $7.25 \times 10^{-16}$  (**g**).

### A flexible planning mechanism

If trajectory events reflect behavioural planning generally, they might also have depicted future behaviours when the animal did not proceed immediately to the Home location. Indeed, away-events closely matched the rat's future path regardless of whether the rat's future path took it to the Home location or elsewhere in the arena (Fig. 6a, c). For both cases, trajectories matched the future path more than the past path (Fig. 6b, d). We proposed that if trajectory events reflected an active process that could switch between goals, then before non-Home-seeking behaviours, not only would the representation of the non-Home-seeking path be enhanced, but the representation of the Home well would be reduced. Indeed, we found reduced Home representation in non-Home-seeking away-events compared to Home-seeking away-events (Fig. 6e).

We finally proposed that a flexible planning mechanism should be able to specify paths of novel importance (a novel combination of start and end points) over familiar terrain. The animals' behaviour showed evidence of this ability over the first 19 trials of each day. We therefore examined trajectory events during this period of each session. Away-events during this novel period also bore a strong match to the rat's future path (Fig. 6f and Supplementary Videos 3–7), and were closer to the rat's future path than its past path (Fig. 6g).



**Figure 6 | Goal switching and flexibility in trajectory events.** **a, b,** As Fig. 5c (left) and 5d, for away-events preceding behaviour ending at or crossing Home (future path, TH-FP; past path, TH-PP). **c, d,** As **a, b,** for away-events preceding behaviours directed elsewhere (future path, NH-FP; past path, NH-PP). **e,** Mean posterior probability representation of Home for same division of away-events (to Home, TH; not to Home, NH). **f, g,** As Fig. 5c, d, for away-events from the first 19 trials of each session (future path, A19-FP; past path, A19-PP). *P*-values (Wilcoxon rank-sum test):  $4.96 \times 10^{-22}$  (**b**);  $1.12 \times 10^{-13}$  (**d**);  $9.60 \times 10^{-3}$  (**e**);

## Discussion

We have demonstrated that hippocampal SWR-associated trajectory events predict immediate future navigational behaviour. This finding follows a succession of results<sup>8</sup> reporting that SWR-associated sequences occur robustly during the awake state<sup>22–26</sup>, that sequences are not always facsimiles of previous behavioural episodes<sup>22–24,26,33</sup> and can even depict novel combinations of previous experiences<sup>26</sup>, and that sequences can be selective to the extent of not always reporting the most recent experience<sup>26</sup>, or even necessarily experiences from the current environment<sup>25</sup>. Moreover, disruption studies using electrical stimulation contingent on SWR detection have revealed a role for sleep SWRs in learning<sup>34,35</sup>, and a specific role for awake SWRs in working memory but not reference memory<sup>36</sup>, which accords with the flexibility of trajectory events in response to a daily changing goal location<sup>27,30,31</sup>. Regarding our observation of stronger prediction before goal-finding than random foraging, it is likely that during the latter behaviour, an animal repeatedly makes online changes to his planned navigational trajectory, which would reduce its initial predictability. This strategic variability may be reflected in the over-dispersion of place-cell firing rates during random foraging<sup>28,37</sup>. Regarding the mechanism generating trajectory events, low-level mechanisms might have contributed, such as the spatial distributions of place cells' firing rates, although these did not account for the precise depiction of the goal location. Alternatively, it is equally possible that the spatial distributions of firing rates emerged as a consequence of the trajectory events. Simple models of encoding routes via direct experience cannot easily explain either the trial-by-trial switching of trajectory events between different goals (Home-seeking versus non-Home-seeking), or the trajectory events corresponding to novel Random–Home combinations<sup>6,38,39</sup>, although the incorporation of contextual coding for the goal might account for some of this functionality<sup>5,40</sup>. It remains unknown whether trajectory events can reflect the calculation of optimal paths in more challenging navigational tasks that incorporate barriers to movement<sup>41,42</sup>. Finally, we might speculate on how the planning function of trajectory events operates. Trajectory depiction by place cells before behaviour might support a plasticity mechanism

that reinforces the particular path, in a way that can be accessed locally during behaviour<sup>43</sup>. For example, trajectory events might drive associations between places en route and estimates of value<sup>19,31,44,45</sup> or chosen action<sup>44,46</sup> that could be accessed subsequently by local place-cell activation during goal-directed behaviour, perhaps in combination with a local look-ahead mechanism such as theta sequences.

In summary, our data reveal a flexible, goal-directed mechanism for the manipulation of previously acquired memories, in which behavioural trajectories to a remembered goal are depicted in the brain immediately before movement. Such findings address longstanding questions about the role of place cells in navigational learning and planning, as well as broader questions regarding the recall and use of stored memory. In particular, trajectory events relate to hippocampal function in multiple conceptual contexts: as a cognitive map in which routes to goals might be explored flexibly before behaviour<sup>1</sup>, as an episodic memory system engaging in what has been termed 'mental time travel'<sup>47</sup>, and as a substrate for the recall of imaginary events<sup>48,49</sup>. These conceptualizations reflect a continuity with earlier speculations on animals' capacities for inference<sup>50</sup>. Trajectory events offer a new experimental model for the study of these varied functions.

## METHODS SUMMARY

A microdrive array containing 40 independently adjustable, gold-plated tetrodes aimed at area CA1 of dorsal hippocampus (20 tetrodes per hemisphere; 4.00 mm posterior and 2.85 mm lateral to bregma) was implanted in four rat subjects. Final tetrode placement and unit recording were as previously described<sup>22</sup>.

Position information was binned into 2-cm bins. Tuning curves were calculated as the smoothed histogram of firing activity normalized by the time spent per bin. Population events were defined as peaks in a smoothed spike density histogram greater than the mean + 3 standard deviations, bounded by crossings of the mean.

Probability-based decoding of position information from spike trains was performed as previously described<sup>23</sup>, using a time window of 20 ms. Each candidate event was truncated to the longest sequence of time frames in which the peak posterior probability was less than 20 cm from that of the previous frame. Events with fewer than 10 steps in the final sequence or a start-to-end distance less than 40 cm were eliminated from further analysis.

**Full Methods** and any associated references are available in the online version of the paper.

Received 28 September 2012; accepted 21 March 2013.

Published online 17 April 2013.

- O'Keefe, J. & Nadel, L. *The Hippocampus As A Cognitive Map*. (Clarendon, 1978).
- Morris, R. G., Garrud, P., Rawlins, J. N. & O'Keefe, J. Place navigation impaired in rats with hippocampal lesions. *Nature* **297**, 681–683 (1982).
- Scoville, W. B. & Milner, B. Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry* **20**, 11–21 (1957).
- Olton, D. S. & Samuelson, R. J. Remembrance of places past: spatial memory in rats. *J. Exp. Psychol. Anim. Behav. Process.* **2**, 97–116 (1976).
- Levy, W. B. A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus* **6**, 579–590 (1996).
- Redish, A. D. & Touretzky, D. S. The role of the hippocampus in solving the Morris water maze. *Neural Comput.* **10**, 73–111 (1998).
- Koene, R. A., Gorchetchnikov, A., Cannon, R. C. & Hasselmo, M. E. Modeling goal-directed spatial navigation in the rat based on physiological data from the hippocampal formation. *Neural Netw.* **16**, 577–584 (2003).
- Foster, D. J. & Knierim, J. J. Sequence learning and the role of the hippocampus in rodent navigation. *Curr. Opin. Neurobiol.* **22**, 294–300 (2012).
- Hok, V. et al. Goal-related activity in hippocampal place cells. *J. Neurosci.* **27**, 472–482 (2007).
- Wood, E. R., Dudchenko, P. A., Robitsek, R. J. & Eichenbaum, H. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* **27**, 623–633 (2000).
- Ferbinteanu, J. & Shapiro, M. L. Prospective and retrospective memory coding in the hippocampus. *Neuron* **40**, 1227–1239 (2003).
- Wilson, M. A. & McNaughton, B. L. Dynamics of the hippocampal ensemble code for space. *Science* **261**, 1055–1058 (1993).
- O'Keefe, J. & Recce, M. L. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* **3**, 317–330 (1993).
- Muller, R. U. & Kubie, J. L. The firing of hippocampal place cells predicts the future position of freely moving rats. *J. Neurosci.* **9**, 4101–4110 (1989).
- Skaggs, W. E., McNaughton, B. L., Wilson, M. A. & Barnes, C. A. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**, 149–172 (1996).

16. Jensen, O. & Lisman, J. E. Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. *Learn. Mem.* **3**, 279–287 (1996).
17. Foster, D. J. & Wilson, M. A. Hippocampal theta sequences. *Hippocampus* **17**, 1093–1099 (2007).
18. Johnson, A. & Redish, A. D. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.* **27**, 12176–12189 (2007).
19. Johnson, A., van der Meer, M. A. & Redish, A. D. Integrating hippocampus and striatum in decision-making. *Curr. Opin. Neurobiol.* **17**, 692–697 (2007).
20. Louie, K. & Wilson, M. A. Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* **29**, 145–156 (2001).
21. Lee, A. K. & Wilson, M. A. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* **36**, 1183–1194 (2002).
22. Foster, D. J. & Wilson, M. A. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**, 680–683 (2006).
23. Davidson, T. J., Kloosterman, F. & Wilson, M. A. Hippocampal replay of extended experience. *Neuron* **63**, 497–507 (2009).
24. Diba, K. & Buzsáki, G. Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neurosci.* **10**, 1241–1242 (2007).
25. Karlsson, M. P. & Frank, L. M. Awake replay of remote experiences in the hippocampus. *Nature Neurosci.* **12**, 913–918 (2009).
26. Gupta, A. S., van der Meer, M. A., Touretzky, D. S. & Redish, A. D. Hippocampal replay is not a simple function of experience. *Neuron* **65**, 695–705 (2010).
27. Steele, R. J. & Morris, R. G. Delay-dependent impairment of a matching-to-place task with chronic and intrahippocampal infusion of the NMDA-antagonist D-AP5. *Hippocampus* **9**, 118–136 (1999).
28. Olypher, A. V., Lansky, P. & Fenton, A. A. Properties of the extra-positional signal in hippocampal place cell discharge derived from the overdispersion in location-specific firing. *Neuroscience* **111**, 553–566 (2002).
29. Kentros, C. G., Agnihotri, N. T., Streater, S., Hawkins, R. D. & Kandel, E. R. Increased attention to spatial context increases both place field stability and spatial memory. *Neuron* **42**, 283–295 (2004).
30. Eichenbaum, H., Otto, T. & Cohen, N. J. The hippocampus—what does it do? *Behav. Neural Biol.* **57**, 2–36 (1992).
31. Foster, D. J., Morris, R. G. & Dayan, P. A model of hippocampally dependent navigation, using the temporal difference learning rule. *Hippocampus* **10**, 1–16 (2000).
32. Schmitzer-Torbert, N., Jackson, J., Henze, D., Harris, K. & Redish, A. D. Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience* **131**, 1–11 (2005).
33. Csicsvari, J., O'Neill, J., Allen, K. & Senior, T. 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 (2007).
34. Girardeau, G., Benchenane, K., Wiener, S. I., Buzsáki, G. & Zugaro, M. B. Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neurosci.* **12**, 1222–1223 (2009).
35. Ego-Stengel, V. & Wilson, M. A. Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus* **20**, 1–10 (2010).
36. Jadhav, S. P., Kemere, C., German, P. W. & Frank, L. M. Awake hippocampal sharp-wave ripples support spatial memory. *Science* **336**, 1454–1458 (2012).
37. Jackson, J. & Redish, A. D. Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. *Hippocampus* **17**, 1209–1229 (2007).
38. Buzsáki, G. Two-stage model of memory trace formation: a role for “noisy” brain states. *Neuroscience* **31**, 551–570 (1989).
39. Mehta, M. R., Lee, A. K. & Wilson, M. A. Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* **417**, 741–746 (2002).
40. Gerstner, W. & Abbott, L. F. Learning navigational maps through potentiation and modulation of hippocampal place cells. *J. Comput. Neurosci.* **4**, 79–94 (1997).
41. Poucet, B., Thinusblanc, C. & Chapuis, N. Route planning in cats, in relation to the visibility of the goal. *Anim. Behav.* **31**, 594–599 (1983).
42. Foster, D. & Dayan, P. Structure in the space of value functions. *Mach. Learn.* **49**, 325–346 (2002).
43. Sutton, R. S. in *Neural Networks for Control* (eds Miller, T., Sutton, R. S. & Werbos, P.) Ch. 8 (MIT Press, 1990).
44. Montague, P. R., Dayan, P. & Sejnowski, T. J. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J. Neurosci.* **16**, 1936–1947 (1996).
45. Lansink, C. S., Goltstein, P. M., Lankelma, J. V., McNaughton, B. L. & Pennartz, C. M. Hippocampus leads ventral striatum in replay of place-reward information. *PLoS Biol.* **7**, e1000173 (2009).
46. van der Meer, M. A., Johnson, A., Schmitzer-Torbert, N. C. & Redish, A. D. Triple dissociation of information processing in dorsal striatum, ventral striatum, and hippocampus on a learned spatial decision task. *Neuron* **67**, 25–32 (2010).
47. Tulving, E. Episodic memory: from mind to brain. *Annu. Rev. Psychol.* **53**, 1–25 (2002).
48. Hassabis, D., Kumaran, D., Vann, S. D. & Maguire, E. A. Patients with hippocampal amnesia cannot imagine new experiences. *Proc. Natl Acad. Sci. USA* **104**, 1726–1731 (2007).
49. Buckner, R. L. The role of the hippocampus in prediction and imagination. *Annu. Rev. Psychol.* **61**, 27–48 (2010).
50. Tolman, E. C. *Purposive Behavior in Animals and Men* (The Century Co., 1932).

**Supplementary Information** is available in the online version of the paper.

**Acknowledgements** This work was supported by the Alfred P. Sloan Foundation, The Brain and Behavior Research Foundation (NARSAD Young Investigator Grant) and the National Institutes of Health grant MH085823.

**Author Contributions** B.E.P. and D.J.F. designed the experiment and analyses, B.E.P. collected the data, B.E.P. and D.J.F. wrote the paper.

**Author Information** Reprints and permissions information is available at [www.nature.com/reprints](http://www.nature.com/reprints). The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to D.J.F. ([david.foster@jhu.edu](mailto:david.foster@jhu.edu)).

## METHODS

**Behaviour and data acquisition.** All procedures were approved by the Johns Hopkins University Animal Care and Use Committee and followed US National Institutes of Health animal use guidelines. Behavioural training and in-session recording took place from late afternoon to early evening (rats were housed on a standard, non-inverted, 12-h light cycle).

Adult male Long-Evans rats (10–20 weeks old, 450–550 g) were handled daily and food-restricted to 85–90% of their free-feeding weight and then trained to traverse a 1.8-m linear track to receive a liquid chocolate-flavoured reward (200  $\mu$ l, Carnation) at either end. Rats were trained for the briefer of 20 min or 20 complete laps once per day for at least 10 consecutive days. Linear track training occurred in a room separate and visually distinct from the recording room.

After a rat achieved criterion performance on the linear track (three consecutive days with 20 laps in under 20 min), training on the open field was initiated in a 2 m  $\times$  2 m black arena with 30-cm-high walls and 36 identical, evenly spaced, 1.5-cm-diameter, 3-mm-deep conical reward delivery wells embedded into the floor such that the rim of each well was level with the floor (Fig. 1a). Each well was attached to a tubing system that ran beneath the environment, which allowed any well to be independently and soundlessly filled or emptied by the experimenter via a hand-held syringe. During the filling of a well, no obvious visible or audible cue was available to the rat signifying that a well had been filled. When active, wells were filled with 300  $\mu$ l of chocolate milk. Open-field training took place in the recording room with all room and environmental cues positioned as they would be during the eventual in-session recording.

Open-field training proceeded in four stages. First, each rat underwent one 30-min-long session per day for 2 days in which every available well was filled (and immediately refilled following consumption) and food crumbs were scattered throughout the arena to encourage initial exploration. This was the only stage of training in which non-liquid food was present in the arena. In the second stage of training (3 days), each 30-min-long session began with four filled wells, one per quadrant of the arena. When the reward in one quadrant was consumed, another random well in that quadrant was filled, but only after the rat had left the quadrant and consumed reward from another quadrant. In the third stage (3 days), the final experimental procedure (see below) was begun except that on the interleaved Random trials, two randomly selected wells were filled to make the task easier to complete. When one Random well was discovered and consumed, the second was immediately emptied and the Home well was filled. Finally, on the fourth stage, the rats were trained on the final experimental protocol for the lesser of thirty minutes or for 30 trials until they reached criterion performance (30 trials in less than 30 min for three consecutive days). Every session began by placing the rat in one corner of the arena and then allowing free exploration.

In the final experimental protocol, the Home well was initially filled and was the only filled well in the arena at the start of the session. Once the rat discovered and consumed the Home well reward, a randomly selected well was filled. Only after the rat discovered and consumed the Random well reward was the Home well again filled. A trial consisted of the rat leaving the Home location, discovering and consuming the reward at a Random well and then returning to the Home location and consuming the reward there. At no point in the training were the rats provided with any cue informing them when the Home or a Random well was filled (filling occurred during or immediately after consumption at the prior well). Instead, the rats learned to return to the Home well location without cue after consuming the reward at a filled Random well and to begin searching for a Random well immediately after consuming the reward at Home. The Home well location changed every session, but was constant throughout the session. The location of the Home well on the recording days had never previously been experienced by the rats as a Home well location, although they had sporadically received reward in those locations as Random wells in previous sessions.

After a rat achieved criterion performance on the task, it was surgically implanted with a microdrive array (25–30 g) containing 40 independently adjustable, gold-plated tetrodes aimed at area CA1 of dorsal hippocampus (20 tetrodes in each hemisphere; 4.00 mm posterior and 2.85 mm lateral to bregma). Following surgical implantation, tetrodes were slowly lowered into the CA1 pyramidal layer over the course of 7–10 days. Final tetrode placement and unit recording were as previously described<sup>22</sup>. Each tetrode consisted of a twisted bundle of four 17.8  $\mu$ m platinum/10% iridium wires (Neuralynx), and each wire was electroplated with gold to an impedance of <150 M $\Omega$  before surgery. A bone screw firmly attached to the skull served as ground. During the first 4 or 5 days following implantation, the rat was not re-exposed to the experimental arena. After this recovery time, while tetrodes were still being advanced to the hippocampus, the rat was trained once per day on the final experimental protocol for the lesser of 30 min or 30 trials to familiarize it with navigating the arena with the microdrive and attached wires.

All data were collected using a Neuralynx data acquisition system and an overhead video system that recorded continuously at 60 Hz. The rat's position and head direction were determined via two distinctly coloured, head-mounted LEDs. Analogue neural signals were digitized at 32,556 Hz. Spike threshold crossings (50  $\mu$ V) were recorded at 32,556 Hz. Continuous local field potential data were digitally filtered between 0.1 and 500 Hz and recorded at 3,255.6 Hz. The beginning and end of reward consumption were manually determined from the captured video data.

**Cluster analysis.** Individual units were identified by manual clustering based on spike waveform peak amplitudes using custom software (xclust2, M. A. Wilson). Only well-isolated units were included in the analysis. Modified  $L_{\text{ratio}}$  values<sup>32</sup> were calculated for each cluster to confirm cluster quality using the peak amplitude of each waveform as the feature set. Briefly, the  $L_{\text{ratio}}$  value of cluster  $C$  is

$$L_{\text{ratio}} = \left( \sum_{i \notin C} \left( 1 - \text{CDF}_{\chi_{df}^2} \left( D_{i,C}^2 \right) \right) \right) / n_s$$

where  $n_s$  is the total number of spikes recorded on the tetrode throughout the experiment,  $i \notin C$  is the set of spikes which are not members of cluster  $C$ ,  $D_{i,C}^2$  is the Mahalanobis distance of spike  $i$  from cluster  $C$ , and  $\text{CDF}_{\chi_{df}^2}$  is the cumulative distribution function of the  $\chi^2$  distribution with  $df = 4$ . We modified the original equation for  $L_{\text{ratio}}$  to allow for comparison between tetrodes with different numbers of spikes and between experiments of varying time spans. As the original equation is a sum, even well-isolated clusters will necessarily have larger  $L_{\text{ratio}}$  values for particularly long experimental sessions or if they occur on tetrodes with large numbers of spikes. Thus, we normalized the sum by the total number of spikes recorded on the tetrode.

Clustered units that may correspond to putative inhibitory neurons were excluded on the basis of spike width and mean firing rate. To ensure accurate decoding of hippocampal events, only rats in which we obtained at least 100 simultaneously recorded place units were used for subsequent analysis.

**Decoding spatial location.** Position was binned (2 cm) and position tuning curves (place fields) were calculated as the smoothed (Gaussian kernel, standard deviation of 4 cm) histogram of firing activity normalized by the time spent per bin. Only periods of time when the rat was moving faster than 5 cm  $s^{-1}$  were used to determine place fields. Units were considered to have a place field if the unit was classified as excitatory and the peak of the tuning curve was >1 Hz.

A memoryless probability-based decoding algorithm<sup>23</sup> was used to estimate the rat's position throughout the experiment based on the unit position tuning curves and the spike trains. Briefly, the probability of the animal's position (pos) across  $M$  total position bins given a time window ( $\tau$ ) containing neural spiking (spikes) is

$$\Pr(\text{pos}|\text{spikes}) = \bigcup_{j=1}^M \bigcup$$

where

$$\bigcup = \left( \prod_{i=1}^N f_i(\text{pos})^{n_i} \right) e^{-\tau \sum_{i=1}^N f_i(\text{pos})}$$

and  $f_i(\text{pos})$  is the position tuning curve of the  $i$ -th unit, assuming independent rates and Poisson firing statistics for all  $N$  units and a uniform prior over position. A time window of 250 ms was used to estimate the rat's position on a behavioural timescale. A time window of 20 ms was used to estimate position during candidate population events.

**Sequential event analysis.** A histogram (1-ms bins) of all clustered units for times when the rat's velocity was less than 5 cm  $s^{-1}$  was smoothed (Gaussian kernel, standard deviation of 10 ms). Population events were defined as peaks in the smoothed histogram greater than the mean + 3 standard deviations. Start and end boundaries for each population event were defined as the points where the smoothed histogram crossed the mean. To prevent estimation artefacts, the time window boundaries for each candidate event were adjusted inward (if necessary) to ensure that the first and last estimation bins contained a minimum of 2 spikes. Candidate events in which fewer than 10% of the clustered units participated or with boundaries less than 50 ms or greater than 2,000 ms apart were excluded from analysis.

For each candidate event, the rat's position was estimated using the probability-based decoding algorithm described above with a 20-ms time window, advanced in 5-ms increments throughout the putative event. Following position estimation, each candidate replay event was truncated to the longest sequence of time frames with peak posterior probability less than 20 cm from that of the previous frame. Candidate events with fewer than 10 steps in the final sequence or a start-to-end distance less than 40 cm were eliminated from future analysis. The remaining candidate events were categorized as 'trajectory events'.

For trajectory event quantification, the posterior probabilities for every time frame of each trajectory event were summed across time. For comparison between away-events and home-events, these sums were normalized for the number of time-frames in each event. For all analyses requiring per-well quantification, the arena was subdivided by drawing an imaginary line equidistant between each well, resulting in 36 regions, each encompassing an approximately  $33 \times 33$  cm area (Supplementary Fig. 4). Quantification for all event trajectory analysis in which the rat's location was not specifically examined did not include the area within 15 cm of the rat's physical location at the time of the event to avoid initiation bias.

For all trajectory events, a Monte-Carlo *P*-value was calculated using two shuffle methods: randomly shuffling cell identity and randomly shuffling each cell's place field in both the *x* and *y* dimensions. The *P*-value was calculated as  $(n + 1)/(r + 1)$ , where *n* is the number of shuffles that met the criteria to be classified as a trajectory event and *r* is the total number of shuffles. 5,000 shuffles were used for both methods. All candidate events that met our criteria to be classified as trajectory events had a *P*-value less than 0.02 for both shuffle methods.

To quantify the precise spatial correlation between trajectory events and the rat's future/past path, each trajectory event was transformed into a vector of the peak posterior probabilities for each time frame of the event. Using the rat's physical location at the time of the event as the centre, concentric rings were drawn around the rat with radial increments of 2 cm, starting with a radius of 15 cm. For each ring, the first crossing for the event vector and the rat's future or past path were determined and the angular displacement (the minor arc along the

ring's circumference, normalized by the ring's radius) was calculated between these points. This value was compared to that obtained from 2,000 randomly selected events (chosen from across all sessions) which were spatially relocated so that the rat's physical location at the time of the random event matched the rat's physical location at the time of the trajectory event to generate a Monte-Carlo *P*-value.

**Local field potential analysis.** For each tetrode, one representative electrode was selected and the local field potential signal was analysed. To examine SWRs, the local field potential was band-pass filtered between 150 and 250 Hz, and the absolute value of the Hilbert transform of this filtered signal was then smoothed (Gaussian kernel, s.d. = 12.5 ms). This processed signal was averaged across all tetrodes and ripple events were identified as local peaks with an amplitude greater than 3 s.d. above the mean, using only periods when the rat's velocity was less than  $5 \text{ cm s}^{-1}$ . The start and end boundaries for each event were defined as the point when the signal crossed the mean. For theta-band power analysis, the raw local field potential trace was band-pass filtered between 4 and 12 Hz and the absolute value of the Hilbert transform of the filtered signal was calculated. The *z*-score theta power for each electrode was determined for every time point of the 60 Hz position data and for 100–200 ms before and after each identified trajectory event. For power spectral density analysis, 100 ms non-overlapping temporal bins were used to compute the spectrograms. A *z*-score was calculated for each frequency band across the entire behavioural session. The SWR or trajectory event triggered spectrograms use the peak of the ripple power or the peak of the spike density, respectively, as time zero.