LETTERS

Reverse replay of behavioural sequences in hippocampal place cells during the awake state

David J. Foster[1] & Matthew A. Wilson[1]

The hippocampus has long been known to be involved in spatial navigational learning in rodents[1,2], and in memory for events in rodents[3,4], primates[5] and humans[6]. A unifying property of both navigation and event memory is a requirement for dealing with temporally sequenced information. Reactivation of temporally sequenced memories for previous behavioural experiences has been reported in sleep in rats[7,8]. Here we report that sequential replay occurs in the rat hippocampus during awake periods immediately after spatial experience. This replay has a unique form, in which recent episodes of spatial experience are replayed in a temporally reversed order. This replay is suggestive of a role in the evaluation of event sequences in the manner of reinforcement learning models. We propose that such replay might constitute a general mechanism of learning and memory.

We used multiple single-unit recording techniques[9] to measure hippocampal neural activity during periods of running and stopping in four rats. Two sessions were recorded per animal, one on a familiar track and one on a new track. During each session, the animal ran several laps, with each lap consisting of running from one end of the track to the other and back again. Within a given lap, the animal stopped at each end to consume food from a food well. After consuming the food, the animal would wait of its own accord in the same position for a short period of time that varied from lap to lap (Fig. 1a). The behaviour of the animal during this time varied between grooming, whisking or being still. The animal would then turn around and immediately begin running again.

For each recording session, we first characterized the activity of neurons in terms of their place fields[10] during locomotion, as measured using the spikes from all laps (with each running direction considered separately; Fig. 1a). Neurons satisfying minimum firing rate and waveform criteria were selected (see Methods), and their place fields were ordered according to the position of the field peaks (Fig. 1c) in order to generate a probe sequence. This probe sequence was then used to examine patterns of activity in cells during individual laps (Fig. 2a, b). While an animal was running, cells fired in order with respect to position, as expected from their place fields. However, during the stopping periods immediately after running, regularly occurring instances of coincident spiking were evident, involving many of the cells in the probe sequence. Notably, within each coincident event, the sequence of cell activation was in reverse order with respect to the probe sequence, and spanned the equivalent of the entire track, on a timescale of hundreds of milliseconds (Fig. 2c).

To quantify the effect for each recording session, we first identified coincident spiking events during stopping periods that involved a large proportion of the cells in the probe sequence for that recording session (see Methods). For each event, the rank-order correlation between cell number and time was calculated, together with a probability[11]. Examples of significant ($P < 0.05$) events are shown in Fig. 3 (Supplementary Figs S2–S5 show example events for all four animals). For each of the eight sessions, over both directions, the distribution of correlation values of all events (regardless of P value) was found to be significantly different from (that is, significantly negative with respect to) the distribution of correlation values of all events with the cell-order parameter shuffled randomly (Fig. 4; P values in figure legend). Hence, the occurrence of reverse replay events was significantly greater than would be expected by chance. The correlation distribution of all events across all four new sessions was significantly different from (that is, more negative than) the distribution of all events across all four familiar sessions (two-tailed Kolmogorov–Smirnov test, $P = 1.13 \times 10^{-10}$), indicating that the phenomenon is more readily observable in a new environment. A number of cells were bidirectional, in that they did not have a peak firing rate in a preferred direction that was at least double that in the opposite direction (52% bidirectional neurons in the new sessions; 35% in the familiar sessions), raising the possibility that apparently reverse replay events merely reflected forward replay of neurons in

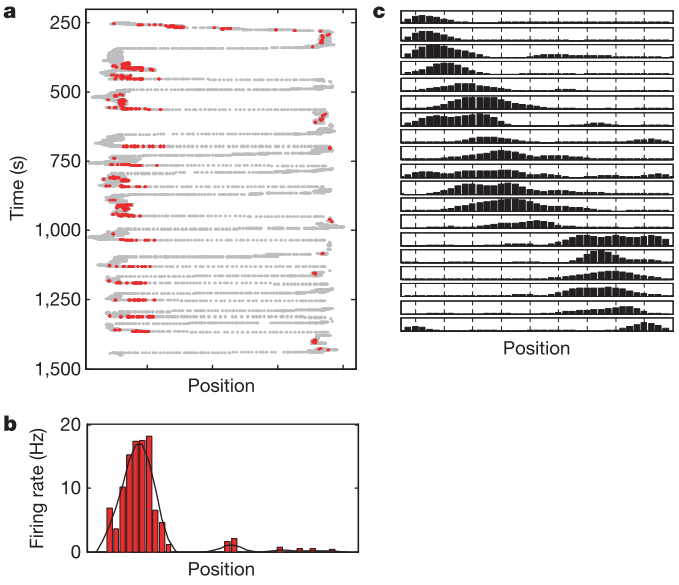

Figure 1 | A sequence of place fields. a, The position of a rat during one recording session is shown in grey as a function of time. Stopping periods at each end could exceed one minute. Spikes emitted by a single hippocampal place cell while the animal faced rightwards are shown in red. **b**, The place field of the single cell shown in Fig. 2a. **c**, Simultaneous recording of 128 cells, of which 26 cells had place fields on the track. Nineteen cells with fields in the rightward direction were ordered by peak to generate a probe sequence.

[1]Picower Institute for Learning and Memory, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 77 Massachusetts Avenue, 46-5223, Cambridge, Massachusetts 02139, USA.

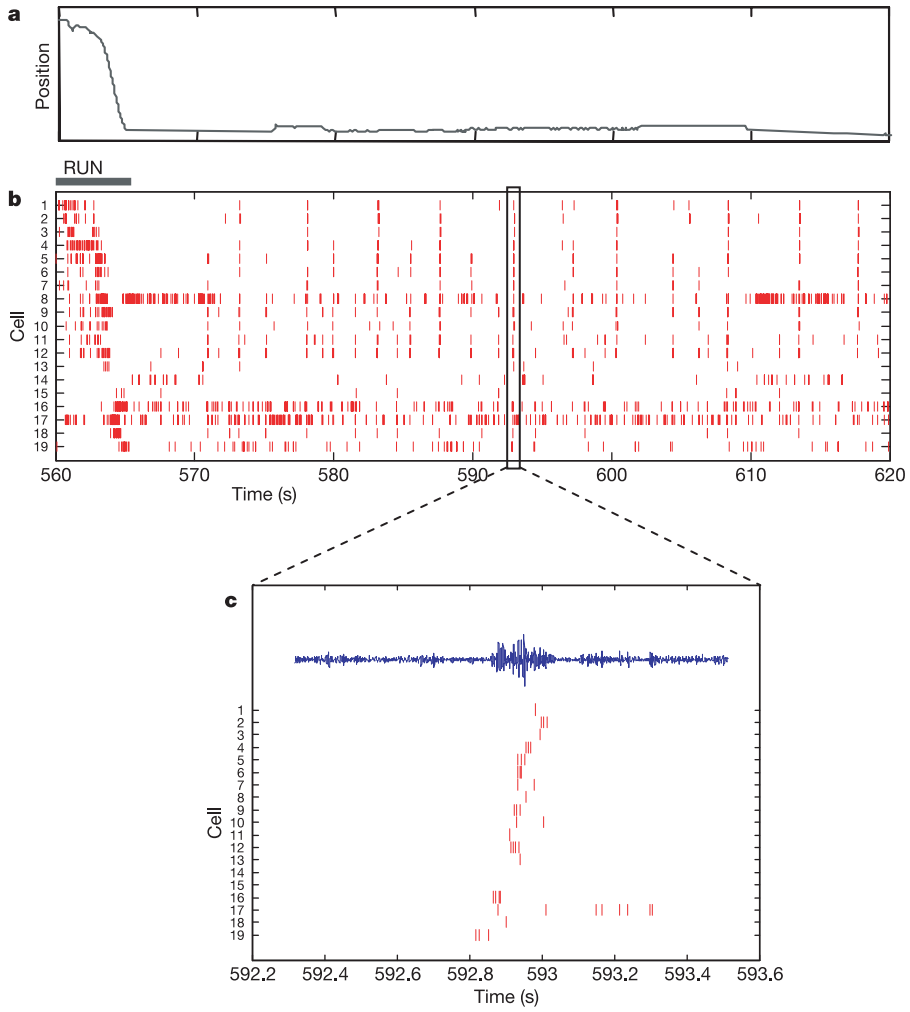

Figure 2 | Reverse replay events during a single lap. **a**, The position of the rat as function of time. **b**, Spikes are shown for each of the cells from Fig. 1c, in the same sequential order. The *x* axis (time) is the same as, and aligned with, panel **a**. During the stopping period, coincident spiking events are visible as narrow, vertical lines. **c**, A section from **b** with the *x* axis expanded to reveal reverse replay. In blue, the simultaneously recorded hippocampal EEG shows a co-occurring ripple event.

the opposite direction. Each of 486 reverse replay events was assessed with the bidirectional cells removed, of which 117 remained significant ($P < 10^{-13}$ under a binomial distribution). Only 8 of the 486 corresponding simultaneous events in the other direction remained significant with the bidirectional cells removed ($P = 1$). Hence, unidirectional cells from the probe sequences for the preferred direction showed significant reverse replay, whereas unidirectional cells for the opposite direction did not show significant forward replay.

Most stopping periods with reverse replay showed multiple reverse replay events (Fig. 4b). Reverse replay occurred even after the first lap on a new track (Fig. 3). Reverse replay events were coincident with ripples in the hippocampal electroencephalogram (EEG; Fig. 4c), which are characteristic of hippocampal activity during both awake, non-running periods and sleep[12–14]. The question remained as to whether reverse replay reflected immediate experience, and so memory for the experiential sequence, or whether the replay could occur in the absence of immediate experience, reflecting some pre-existing expectation of sequential order. In six sessions, we recorded cell activity after the animal had been placed on the track but before running, during which time the animal was still and facing away from the track, hence in a similar physical state to that occupied during subsequent stopping periods in that location. None of these periods showed reverse replay, although the periods ranged between 42.3 s and 424.3 s in duration. A possible model for the generation of reverse replay sequences that encompasses these data is presented in Supplementary Fig. S6.

The hippocampus has long been known to be necessary for learning in sequential decision problems such as navigation[1–3]. Sequential decision problems suffer from the well-known temporal credit assignment problem—that of relating reward information that might occur only at the end of a sequence of events to the individual events within that sequence. A classic solution to this problem is to propagate value information from the rewarded location backwards along incoming trajectories[15–22]. In the brain, reverse replay could be paired with a fast-onset, slowly decaying dopamine signal to learn a representation of value, thus providing a value gradient that the animal could follow during subsequent goal-finding behaviour (Supplementary Fig. S7). Hence, reverse replay in the hippocampus might have a critical role in support of learning in hippocampus-dependent tasks. The finding that reverse replay is more readily observable in a new environment than a familiar one is consistent with such a role.

Reverse replay during the awake state can be contrasted with replay in sharp waves during slow-wave sleep, in which episodes of spatial experience are replayed in the same temporal order as that in which they were experienced[8]. This re-expression of events while the animal occupies an entirely different physical and temporal context, as well as a different behavioural state, may have a role in memory consolidation during sleep[23,24]. When awake, reverse replay occurs *in situ*, allowing immediately preceding events to be evaluated in precise temporal relation to a current, anchoring event, and so may be an integral mechanism for learning about recent events. Moreover, by converting single experiences into multiple reverse events, even after the first encounter in a new environment, awake replay represents efficient use of hard-won experience. Understanding this replay is likely to be critical to understanding how animals learn from experience.

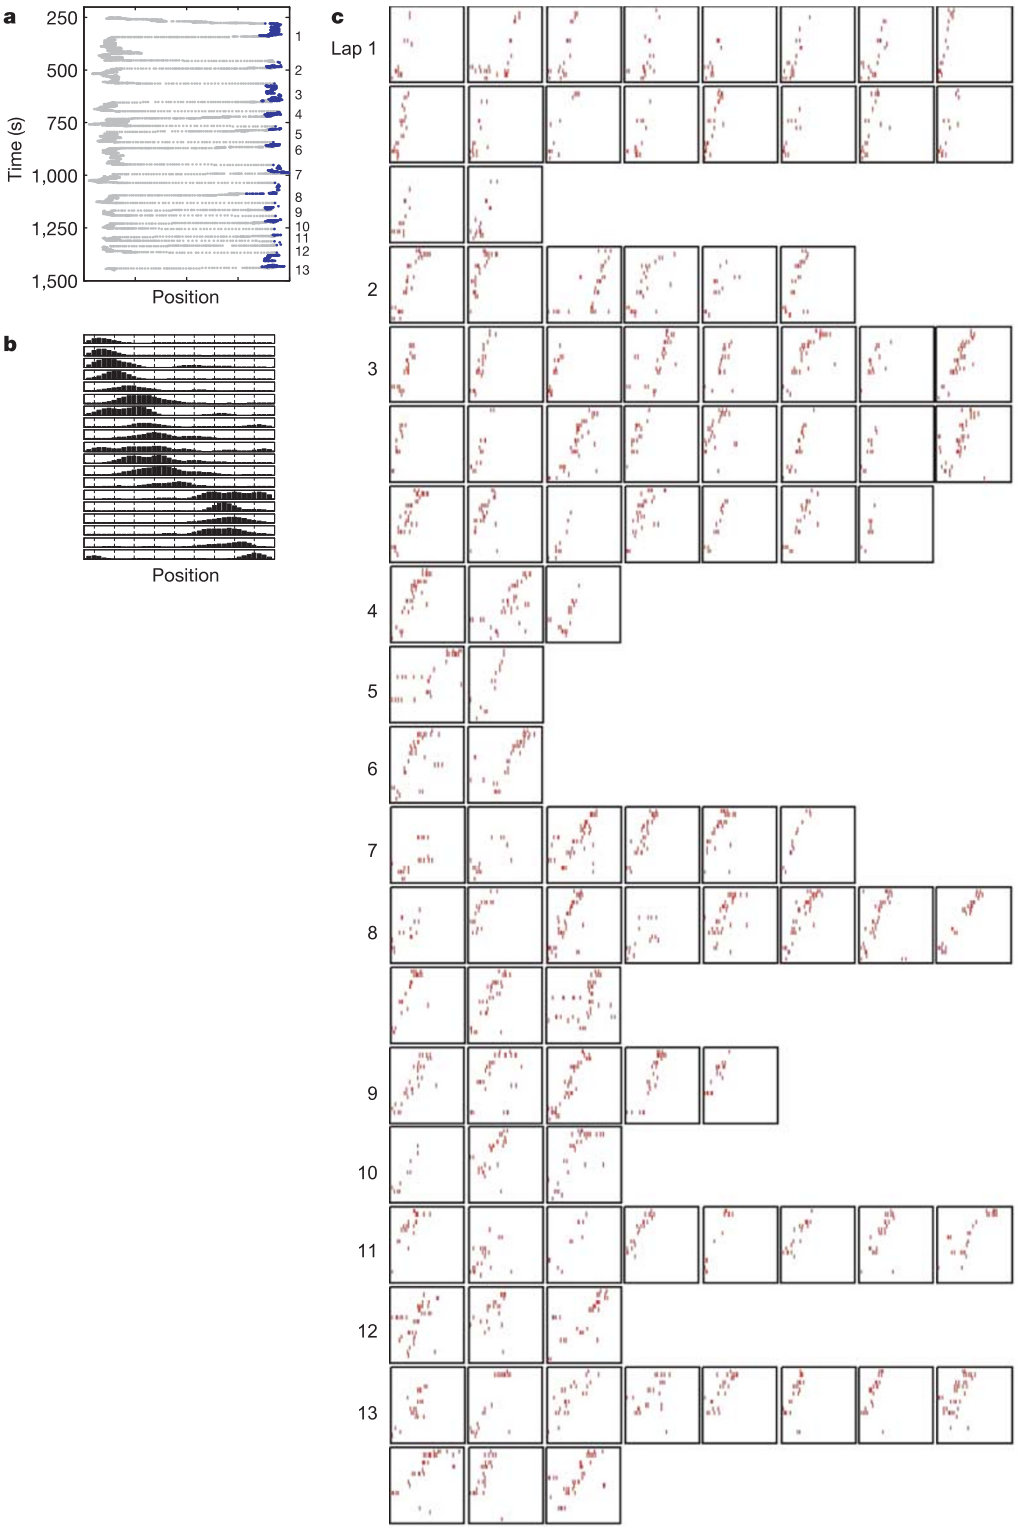

Figure 3 | Reverse replay events, by lap, for an entire recording session. **a**, The session follows rat 2 on the new track. Position during running periods is shown in grey; stopping periods (for the rightward direction only) are shown in blue. Laps are numbered on the right. **b**, The probe sequence for the rightward direction, showing the position of the 19 place fields in the sequence. **c**, Reverse replay events by lap. The *y* axis in each case is cells 1–19. The *x* axis in each case is a fixed time-window of 288 ms.

METHODS

Electrophysiology and behavioural apparatus. In each of four rats, a multiple electrode microdrive array[9] consisting of either 18 (rats 1 and 4) or 17 (rats 2 and 3) independently adjustable tetrodes was implanted above the right dorsal hippocampus (4 mm posterior, 2.2 mm lateral with respect to bregma), and the tetrodes were lowered over the course of several days until they rested in the CA1 pyramidal cell layer. The remaining details of the procedure were as previously described[8]. Direction was measured by the relative position of two tracker diodes mounted to the front and rear of the tetrode drive. Linear tracks (162 cm long) were used for both new and familiar sessions for rats 1–3, and a U-shaped track (205 cm long, 45 cm wide) was used for both new and familiar sessions for rat 4.

Place-field analysis. Position was linearized for each session to yield a scalar value of distance along the track. A histogram of spikes from each cell was calculated over position bins and was normalized by the time spent by the animal in each bin, to yield a place field. Fields were velocity-filtered to exclude times when the speed of the animal was below 5.4 cm s^{-1}. In order to assign a peak value, the histogram was smoothed (as shown by the black line in Fig. 2b). Cells with a peak firing rate of at least 5 Hz were included in the probe sequence, with the exception of putative inhibitory interneurons, which were identified as cells with a mean peak-to-trough spike width of less than 0.35 ms.

Spike-train analysis. A spike train was constituted from all spikes (from all cells in the probe sequence) that occurred during stopping periods while the animal faced in the direction in which it had just run. This spike train was then broken between every pair of successive spikes separated by more than 50 ms, to form a

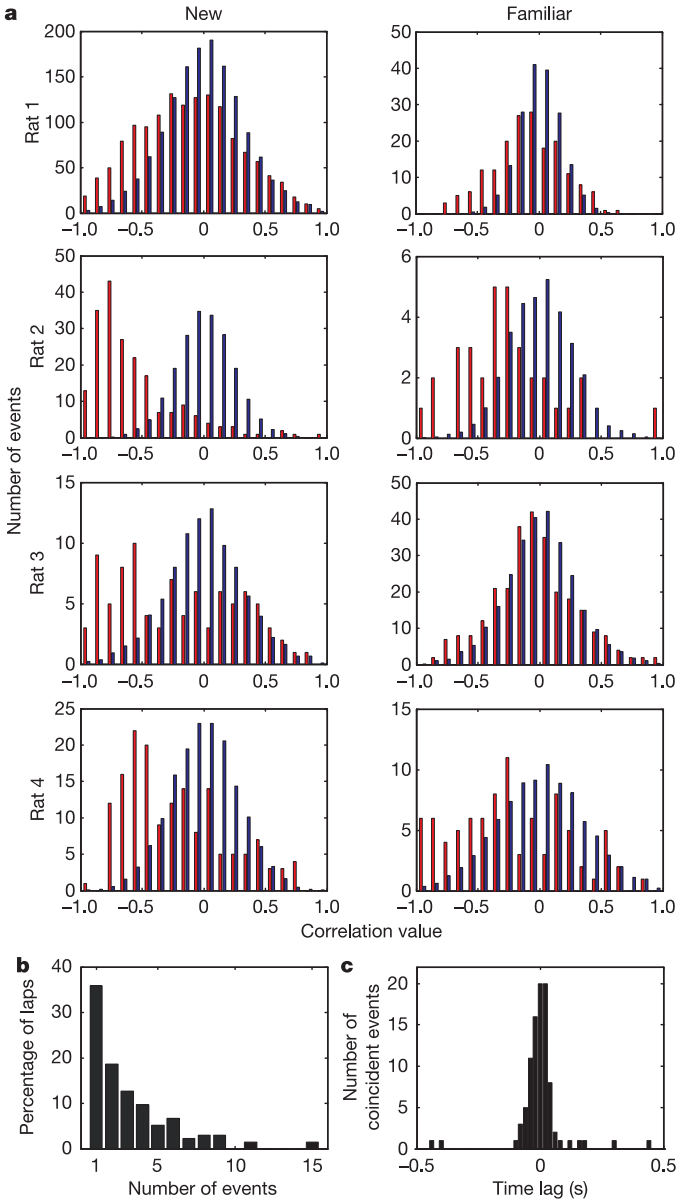

Figure 4 | Analysis of reverse replay across all recording sessions. a, For each session, a histogram of the rank-order correlation values of all events is shown in red, and a histogram of shuffled correlation values is shown in blue (see Methods). The two distributions were statistically different. *P* values for each session were as follows, where *n* is the total number of events. New: Rat 1, $n = 1{,}425$, $P = 3.04 \times 10^{-41}$; Rat 2, $n = 202$, $P = 4.00 \times 10^{-98}$; Rat 3, $n = 91$, $P = 6.52 \times 10^{-10}$; Rat 4, $n = 160$, $P = 4.88 \times 10^{-21}$. Familiar: Rat 1, $n = 178$, $P = 6.69 \times 10^{-8}$; Rat 2, $n = 33$, $P = 4.05 \times 10^{-6}$; Rat 3, $n = 275$, $P = 0.0067$; Rat 4, $n = 88$; $P = 1.32 \times 10^{-8}$. The percentage of events with significant reverse correlations was as follows, by session: New, Rat 1, 13%; Rat 2, 72%; Rat 3, 31%; Rat 4, 29%; Familiar, Rat 1, 19%; Rat 2, 30%; Rat 3, 6%; Rat 4, 16%. **b**, Histogram of the number of significant reverse events per stopping period, for those stopping periods with at least one significant event. **c**, Cross-correlogram of significant reverse replay events with hippocampal sharp waves, for an example session in which there were 94 coincident events out of a total of 146 replay events.

large set of proto-events. Those proto-events in which at least one-third of the cells in the probe sequence fired at least one spike were then selected as events. The few events longer than 500 ms in duration were rejected as a potential source of spurious correlations. For each event, 100 shuffled events were created by randomly permuting the cell-order parameter. The histograms in Fig. 4a were normalized (by dividing by 100) to allow visual comparison with the original distributions. A non-parametric, two-sample Kolmogorov–Smirnov test was used to determine whether the distributions were significantly different.

Ripple identification. Sharp waves reverse at about the electrode depth corresponding to maximum cell yield in the hippocampus, making it difficult to measure sharp waves directly. However, they co-occur with transient, high-frequency events called ripples (100–400 Hz). Ripples were identified as reported previously[8]. A single time of occurrence for each ripple was calculated as the mean of the start and end times of the ripple. A single time was similarly calculated for each replay event. These times were used to generate a cross-correlogram, which was not normalized, so that the *y* axis of Fig. 4c is in numbers of coincident events. Values for total numbers of coincident events cited in the text were found by summing the values of bins between −50 ms and 50 ms.

Received 13 December 2005; accepted 20 January 2006.
Published online 12 February 2006.

1. O'Keefe, J. & Nadel, L. *The Hippocampus as a Cognitive Map* (Clarendon, London, 1978).
2. Morris, R. G. M., Garrud, P., Rawlins, J. N. P. & O'Keefe, J. Place navigation impaired in rats with hippocampal lesions. *Nature* **297,** 681–683 (1982).
3. Olton, D. S. & Samuelson, R. J. Remembrance of places past: spatial memory in rats. *J. Exp. Psychol. Anim. Behav. Process.* **2,** 97–116 (1976).
4. Steele, R. J. & Morris, R. G. M. 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).
5. Gaffan, D. Scene-specific memory for objects: a model of episodic memory impairment in monkeys with fornix transection. *J. Cogn. Neurosci.* **6,** 305–320 (1994).
6. Scoville, W. B. & Milner, B. Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry* **20,** 11–21 (1957).
7. Louie, K. & Wilson, M. A. Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* **29,** 145–156 (2001).
8. Lee, A. K. & Wilson, M. A. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* **36,** 1183–1194 (2002).
9. Wilson, M. A. & McNaughton, B. L. Dynamics of the hippocampal ensemble code for space. *Science* **261,** 1055–1058 (1993).
10. O'Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely moving rat. *Brain Res.* **34,** 171–175 (1971).
11. Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. *Numerical Recipes in C: The Art of Scientific Computing* 2nd edn (Cambridge Univ. Press, Cambridge, 1992).
12. Ranck, J. B. Jr. Studies on single neurons in dorsal hippocampal formation and in septum in unrestrained rats. I. Behavioral correlates and firing repertoires. *Exp. Neurol.* **42,** 461–531 (1973).
13. Buzsaki, G. Hippocampal sharp waves: their origin and significance. *Brain Res.* **398,** 242–252 (1986).
14. Buzsaki, G. A two-stage model of memory trace formation: a role for 'noisy' brain states. *Neuroscience* **31,** 551–570 (1989).
15. Kaelbling, L. P., Littman, M. L. & Moore, A. W. Reinforcement learning: a survey. *J. Artif. Intell. Res.* **4,** 237–285 (1996).
16. Sutton, R. S. Learning to predict by the methods of temporal differences. *Mach. Learn.* **3,** 9–44 (1988).
17. Schultz, W., Dayan, P. & Montague, P. R. A neural subtstrate of prediction and reward. *Science* **275,** 1593–1599 (1997).
18. Foster, D. J., Morris, R. G. M. & Dayan, P. A model of hippocampally dependent navigation, using the temporal difference learning rule. *Hippocampus* **10,** 1–16 (2000).
19. Lin, L. J. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Mach. Learn.* **8,** 293–321 (1992).
20. Sutton, R. S. & Singh, S. P. in *Proceedings of the Eighth Yale Workshop on Adaptive and Learning Systems* 31–36 (Yale Univ., New Haven, 1994).
21. Cichosz, P. An analysis of experience replay in temporal difference learning. *Cybernet. Syst.* **30,** 341–363 (1999).
22. Reynolds, S. *Experience Stack Reinforcement Learning for Off-policy Control*. Cognitive Science Technical Report CSRP-02–1 (Birmingham Univ., Birmingham, 2002).
23. Stickgold, R., Whidbee, D., Schirmer, B., Patel, V. & Hobson, J. A. Visual discrimination task improvement: a multi-step process occurring during sleep. *J. Cogn. Neurosci.* **12,** 246–254 (2000).
24. Siapas, A. G. & Wilson, M. A. Coordinated interactions between hippocampal wripples and cortical spindles during slow-wave sleep. *Neuron* **21,** 1123–1128 (1998).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements We thank V. Ego-Stengel, P. Dayan and M. Fee for comments on the manuscript. This work was supported by an NIH grant to M.A.W.

Author Information Reprints and permissions information is available at npg.nature.com/reprintsandpermissions. The authors declare no competing financial interests. Correspondence and requests for materials should be addressed to D.J.F. (djfoster@mit.edu).