

Theta-paced flickering between place-cell maps in the hippocampus

Karel Jezek^{1,2}, Espen J. Henriksen¹, Alessandro Treves^{1,3}, Edvard I. Moser¹ & May-Britt Moser¹

The ability to recall discrete memories is thought to depend on the formation of attractor states in recurrent neural networks^{1–4}. In such networks, representations can be reactivated reliably from subsets of the cues that were present when the memory was encoded, at the same time as interference from competing representations is minimized. Theoretical studies have pointed to the recurrent CA3 system of the hippocampus as a possible attractor network^{3,4}. Consistent with predictions from these studies, experiments have shown that place representations in CA3 and downstream CA1 tolerate small changes in the configuration of the environment but switch to uncorrelated representations when dissimilarities become larger^{5–9}. However, the kinetics supporting such network transitions, at the subsecond timescale, is poorly understood. Here we show in rats that instantaneous transformation of the spatial context does not change the hippocampal representation all at once but is followed by temporary bistability in the discharge activity of CA3 ensembles. Rather than sliding through a continuum of intermediate activity states, the CA3 network undergoes a short period of competitive flickering between preformed representations of the past and present environment before settling on the latter. Network flickers are extremely fast, often with complete replacement of the active ensemble from one theta cycle to the next. Within individual cycles, segregation is stronger towards the end, when firing starts to decline, pointing to the theta cycle as a temporal unit for expression of attractor states in the hippocampus. Repetition of pattern-completion processes across successive theta cycles may facilitate error correction and enhance discriminative power in the presence of weak and ambiguous input cues.

The place-cell population of the hippocampus is thought to create a neural representation of the spatial environment¹⁰. Accumulating evidence indicates that environments are generally represented in hippocampal cells by many discrete maps, each corresponding to a distinct environment or a unique experience within the environment^{5,9,11–13}. Which map is active at any given time depends on external sensory inputs as well as recent history^{7,14}. Incongruity between the active map and sensory inputs may lead to partial or complete replacement of the active representation^{15,16}. The kinetics of map substitutions has remained elusive owing to a shortage of experimental and analytical tools for subsecond-timescale neural population analyses.

Here, we developed tools to determine how local network activity evolves in the hippocampus in response to sudden changes in the cues that define spatial context. Rats with tetrodes in CA3 (Supplementary Fig. 1) were first trained on separate trials in two boxes with different sets of light cues in a dark room (boxes A and B). The procedure favours the development of uncorrelated place representations in A and B^{5,9} (Supplementary Fig. 2). After several days of training in each box, the rats started in one of the environments; then, after 40–60 s, the cues were switched instantaneously to those of the other environment, effectively ‘teleporting’ the rat from A to B or vice versa. A total of 169 such trials, hereafter referred to as teleportation trials, were performed,

with an average of 33 ± 3 active cells per day (mean \pm s.e.m.; total of 358 active cells; 11 days; 6 rats).

To examine the evolution of network activity after the cue change, we first established expected firing patterns for all locations in each environment. Firing rates were determined for each cell in each spatial bin of each box on separate reference trials (30×30 bins; Fig. 1a). As expected for CA3 place cells¹⁷, the subsets of active cells in the two boxes overlapped minimally: cells with activity in both A and B generally fired at unrelated positions (spatial correlation: 0.112 ± 0.019). The nearly orthogonal nature of the baseline representations made it possible, in principle, to infer from any subsequent cell sample which of the two environments was represented in the hippocampal network at the time of recording. Thus, in the next step, we compared the evolution of activity over successive theta cycles in the teleportation test with activity at corresponding locations in the reference environments. The theta phase with the lowest overall firing rate was used to segment the recording from each teleportation trial into individual theta cycles (Fig. 1b, c and Supplementary Fig. 3). A population vector was then constructed for each theta cycle, consisting of the number of spikes of each of the C simultaneously recorded cells (Fig. 1c). For each theta cycle, the population vector was correlated with vectors of average firing rates (in Hz) for the same location in each reference recording. We first used a simple linear measure of correlation, the un-normalized dot product (DP), between the test population vector \mathbf{r} and each of the two reference vectors, \mathbf{r}_A and \mathbf{r}_B . If \mathbf{r}_A and \mathbf{r}_B were strictly orthogonal and after teleportation test vectors were to be linear combinations $\mathbf{r} \approx c_A \mathbf{r}_A + c_B \mathbf{r}_B$, the DP measure would allow a direct estimate of the coefficients $c_A(t)$ and $c_B(t)$, that is, of the strength with which the two representations are relayed to downstream neurons at time t ^{2,4}. Only trials with continuous theta activity were analysed (149 out of 169 teleportation trials).

The switching of spatial cues was followed by a characteristic pattern of network activity (Fig. 2). Before the transformation, most theta cycles correlated strongly with reference vectors for the first environment (denoted as I, whether A or B) but not for the second (II), as expected (Figs 1d and 2a, b). After cue switching, most theta cycles correlated with the reference vectors for II but the magnitude of the dot product was variable and the network occasionally ‘flickered’ back to strong correlations with I. To investigate the statistics of this flickering, we plotted for each theta cycle the correlation of the momentary population vector with the reference vectors from the same environment (x -axis) and the alternative environment (y -axis) (Fig. 2c). The analysis confirmed that cue switching increased the proportion of cycles correlated with the alternative environment but also identified a subset of cycles that correlated modestly with both environments at the same time. This led us to ask if the instantaneous activity during the transition reflected a simple linear combination of fragments of the A and B representations, as would be expected in the absence of population coherence. We consequently determined whether, in the pooled data, the proportion of theta cycles correlating with both environments at

¹Kavli Institute for Systems Neuroscience and Centre for the Biology of Memory, Norwegian University of Science and Technology, Olav Kyrres gate 9, MTF5, 7489 Trondheim, Norway. ²Department of Physiology of Memory, Institute of Physiology, Videnska 1083, Academy of Sciences of the Czech Republic, 142 20 Prague 4-Krc, Czech Republic. ³Cognitive Neuroscience Sector, SISSA International School for Advanced Studies, via Bonomea 265, 34136 Trieste, Italy.

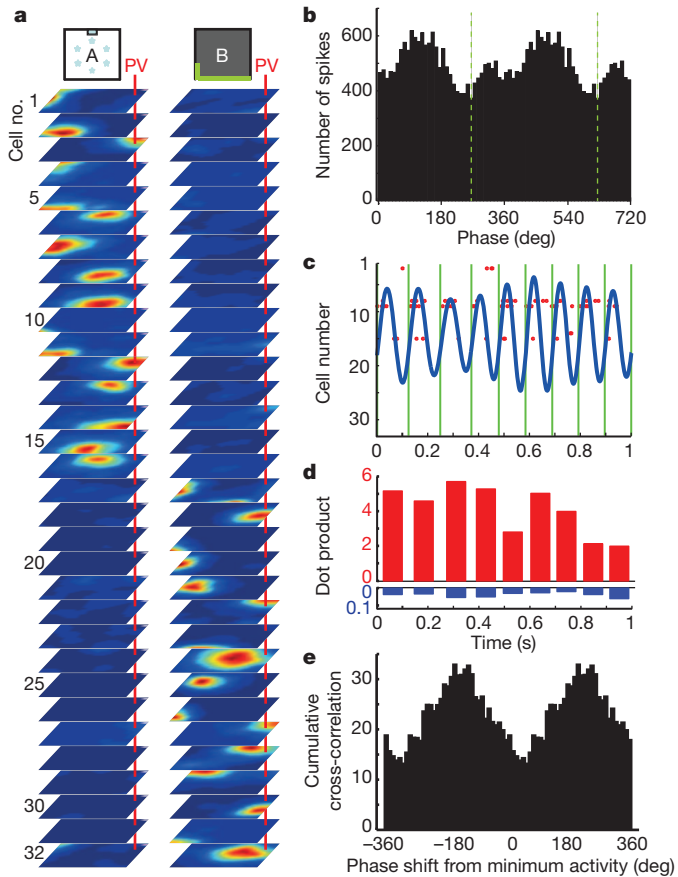


Figure 1 | Procedures for analysing hippocampal transition dynamics.

a, Stack of firing-rate maps in box A (white floor lights; left) and box B (green wall lights; right) for an example set of 32 simultaneously recorded hippocampal CA3 units. Each map shows a colour-coded distribution of firing rates across the square test box (blue, silent; red, maximum). Red line, one of 30×30 population vectors (PV) constructed from the activity of the entire cell ensemble in a given 2×2 cm position bin. Note strong difference in population vectors for A and B. **b**, Theta phase modulation for all pyramidal cells on a representative trial. Spike number is shown as a function of theta phase (bin size 10 deg). Dashed green line, phase with lowest firing rate, used to define boundary between cycles. **c**, Representative spike distribution across theta cycles in the stable state. Rasters of red dots show spike times of individual cells in relation to 6–11 Hz filtered local electroencephalogram (EEG; blue). Green lines indicate theta-cycle boundaries (**b**). The ensemble distribution of activity during one cycle represents the momentary population vector. **d**, Dot-product correlation between momentary population vector and reference vectors at the corresponding position in A (red) and B (blue) during a baseline trial in A. **e**, Cumulative product between correlations with each of the reference environments as a function of the phase for segmentation of theta cycles (0, phase of minimum activity).

the same time ($\mathbf{r} \cdot \mathbf{r}_A > C$, $\mathbf{r} \cdot \mathbf{r}_B > C$, referred to as ‘mixed theta cycles’) was lower than expected if each single unit expressed either one or the other representation independently of the other units. A total of 1.25% of the theta cycles in the recorded data were mixed. This number was lower than in 970 out of 1,000 randomly recombined (shuffled) population vectors (that is, $P < 0.03$; Fig. 2d). The separation between the A- and B-correlated representations was strongest when the cycles were chunked at the point of the lowest average firing rate in the population (Fig. 1b, c, e and Supplementary Fig. 3). Thus, mixed representations existed but were rare. Transitions between orthogonal maps tend to occur in an all-or-none manner, with the entire network flickering coherently at time scales of approximately a tenth of a second.

We subsequently examined the evolution of network activity within the theta cycle. Each cycle was divided into two halves, and mixed states were defined as those half-cycles for which $\mathbf{r} \cdot \mathbf{r}_A > C/2$, $\mathbf{r} \cdot \mathbf{r}_B > C/2$. Mixed population vectors were less abundant than




Figure 2 | Theta cycles correlate with either of the reference environments but rarely with both simultaneously.

a, Top, local hippocampal EEG during spatial cue switching from A to B (filtered at 6–11 Hz). Bottom, dot-product correlation between momentary population vectors and reference vectors from A (red) and B (blue) for successive theta cycles before and after cue switching. Dot products are un-normalized (just divided by the number C of recorded cells). All correlations are positive but, for clarity, A and B correlations are plotted in opposite directions. Green line indicates light switch. EEG and ensemble activity were sampled simultaneously. Note that ensemble activity flickered back to the A representation several times after cue switching. Note also the variation in the dot product. **b**, Another example of network flickering induced by switching of spatial cues. **c**, Matrices reporting the number of cycles falling in each 0.2×10 bin of the dot-product correlations $\mathbf{r} \cdot \mathbf{r}_{A,B}/C$ between momentary population vectors and reference vectors for the present environment (x -axis) or the alternative environment (y -axis). Left, before cue switching. Right, after cue switching (starting from the first cycle correlated with new environment). Note that mixed cycles, defined as cycles with both $\mathbf{r} \cdot \mathbf{r}_A$ and $\mathbf{r} \cdot \mathbf{r}_B$ exceeding C ($x > 1$, $y > 1$; indicated by red lines), were rare. **d**, Histograms showing that the number of mixed states after cue switching (red line) is lower than expected from shuffled versions of the same data ($n = 1,000$; grey histogram). Note that mixed states became less frequent during the second half of the theta cycle.

expected from shuffled data during both half-cycles; however, in the shuffled data, the frequency of recombinations with more mixed population vectors than in the observed data increased from 958/1,000 during the first half ($P < 0.05$) to 1,000/1,000 during the second ($P < 0.001$) (Fig. 2d). The low incidence of mixed population vectors at the end of the theta cycle suggests that representations evolve from partially segregated to fully segregated within each activity period.

We then asked how A-correlated and B-correlated theta cycles were organized in time. Because dot products can vary from 0 to indefinitely large, we switched to Pearson product-moment correlations, which by normalizing the correlations to within a fixed $[-1, +1]$ range allow successive theta cycles to be compared more directly (Fig. 3a–c and Supplementary Figs 4–6). Individual theta cycles were now only considered if at least two cells were active (for higher thresholds, see Supplementary Fig. 7). As observed with the dot products, the

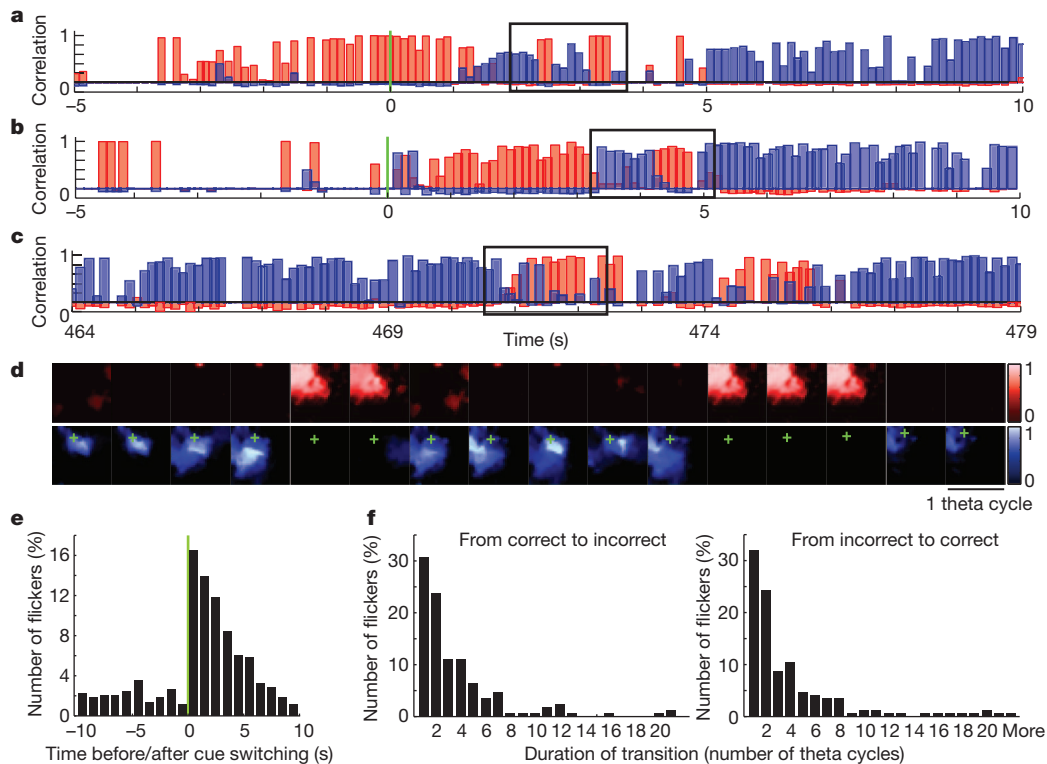


Figure 3 | Temporal dynamics of network flickering. **a–c**, Pearson product-moment correlations showing evolution of population vector correlations after spatial cue switching from box I to II (red, correlation with I; blue, with II; green line, cue switch). Note frequent flickers to the original representation after cue switching in **a** and **b**. **c**, Spontaneous flickering between cue-switch events. Frames: sequences detailed in **d** and Supplementary Fig. 8. **d**, Spatial distribution of correlations between momentary population vectors and reference vectors in the framed area in **a**. Each row shows correlation matrices

for 16 consecutive theta cycles. Top, correlation with I; bottom, with II.

Correlation is colour coded (scale bar). + indicates rat position. **e**, Percentage of flickers to the alternative representation as a function of time before and after cue switching. **f**, Distribution plots showing the time for the network to switch from present ('correct') to past ('incorrect') representation, or vice versa, in number of theta cycles. One cycle means that cycles with alternative representations were consecutive. Note the predominance of immediate transitions.

momentary population vectors correlated strongly with either A or B but rarely with both. Before cue switching, theta cycles were nearly exclusively correlated with reference vectors for the initial environment (I). After cue switching, the network switched almost instantaneously to high correlation with reference vectors for II but then relapsed to I several times during the subsequent seconds (Fig. 3a, b and Supplementary Figs 5 and 6) and occasionally tens of seconds after cue switching (Fig. 3c). These relapses, or flickers, were confined to discrete periods of one or several theta cycles. When the population vector was correlated with reference vectors from other locations in the represented environment, the correlation generally decreased with distance from the animal (Fig. 3d and Supplementary Fig. 8).

To quantify the frequency, timing and duration of flicker events, we defined individual theta cycles as A-correlated if the correlation with reference environment A was above the 95th percentile for $B \times A$ correlations in the reference sessions (that is, more similar to A than 95% of theta cycles in B) and if the correlation with reference environment B was simultaneously below the 5th percentile for $B \times B$ (that is, different from most theta cycles in B) (Supplementary Figs 9 and 10). The analysis showed a clear increase in the frequency of network flickers during the first seconds following the cue switch from I to II, after the network had switched to the II representation for the first time (generally 0–1 s after the cue change; Supplementary Fig. 11 and Supplementary Table 1). The fraction of theta cycles participating in flicker episodes, estimated with Pearson correlations, increased from a stable baseline of 1–3% before cue switching to a level of 10–15% during the first 5 s after the first network switch (Fig. 3e and Supplementary Fig. 12), confirming the tendency revealed by the dot-product analyses (Fig. 2c). The interquartile range of flicker durations increased from 1–1 theta cycles before cue switching to 1–4 during the first 10 s after cue switching (Wilcoxon rank-sum test:

$Z = 2.27$, $P < 0.03$; Supplementary Fig. 11b and Supplementary Table 1). Flicker events were distributed across the entire recording box (Supplementary Fig. 13) and showed no preference for running or heading direction (Supplementary Fig. 14). Flickering was apparent also in CA1 but discrete relapses were clearly less frequent, possibly due to the less reverberating architecture of this subfield (Supplementary Fig. 15).

Transitions between representations occurred within less than a single theta cycle. After the last cycle in a series of non-flicker theta cycles, the alternative representation was fully expressed already in the subsequent theta period in 30.8% of the flicker events where flicker duration could be determined (Fig. 3f). The corresponding percentage of single-cycle returns to the 'correct' representation was 32.0%. In 57.0% of transitions from 'correct' to 'incorrect' reference frame and 62.2% of transitions back from 'incorrect' to 'correct', the time course remained undetermined because the intervening theta cycle contained fewer than two spikes. In each instance of an immediate transition, the network representation was fully developed from the outset, that is, the correlation with the new environment did not increase further within the flicker period ($r = -0.06 \pm 0.11$).

To determine further if the flicker episodes were patterned by the theta oscillation, we finally compared the transition dynamics of our theta-based segmentation procedure with segmentations based on fixed time bins of different width (range 44–500 ms; Supplementary Fig. 3). Among the fixed bins, direct transitions were most abundant when the bins matched the average duration of theta cycles (125 versus 120.4 ms, respectively). The abundance of sharp transitions increased further when the trial was segmented by actual rather than fixed theta periods, with cycles split at the phase with the minimum firing rate (Fig. 1b, c, e and Supplementary Fig. 3). Collectively, these observations suggest strongly that the transitions were paced by the theta rhythm.

Our study provides evidence for competitive interactions between hippocampal representations during changes in spatial reference frame. Although a small subset of the population vectors correlated with both reference environments after the cue change, the number of such mixed states was lower than expected from a sample of independent single units, especially during the second half of each theta cycle. In most cases, the network either switched all at once or flickered between mutually exclusive representations until, after a few seconds, it settled in one of the alternatives. The sustained separation of the neural activity pattern is consistent with the notion that spatial environments, as a whole, are stored as discrete attractors in neural networks of the hippocampus or associated areas such as the entorhinal cortex^{1–4,8,9}. These discontinuities differ from the continuous or quasi-continuous nature of spatial maps for individual environments^{12,18–20}, where sweeps can follow unbroken trajectories, even when retrieval occurs in the absence of actual movement²¹. Changes in attention or experience are likely to generate continuous transitions of the latter type more or less constantly in all environments^{7,22–26}. The low frequency of flickering in the baseline state of the teleportation task suggests, however, that switches to uncorrelated attractor maps are rare and occur primarily when cues are ambiguous or in conflict across sensory modalities.

The time course of flickering episodes has implications for the mechanisms of ensemble activation. It took often only a single theta cycle to fully reactivate a pre-established representation during a flicker event, consistent with models of theta phase precession in which ensembles are activated by propagation through recurrent collaterals after afferent input to a subset of cells early in the theta cycle²⁷. The idea that sensory influences can override attractors at the beginning of the theta cycle, whereas subsequent activity is determined more exclusively by propagation through associative connections, receives further support from the fact that the small number of mixed population vectors occurred primarily during the first half of the theta cycle. The recreation of spatial representations on successive theta cycles in the hippocampus is fundamentally different from the pattern-completion dynamics observed, for example, in the inferior temporal cortex, where the activation may proceed only once per stimulus presentation, at a timescale short enough to facilitate perception but with little opportunity for error correction²⁸. In the hippocampus, repeated convergence to an attractor state might allow the system to self-correct and thereby enhance its discriminative power under conditions where input cues are weak and ambiguous. The present data point to theta cycles as organizational units for this repetitive process but do not preclude additional structure at faster time scales, for example with gamma cycles as units for cell assembly sequences within the theta cycle^{29,30}.

METHODS SUMMARY

Neuronal ensemble activity was recorded from ensembles of CA3 cells while rats foraged in either of two distinct enclosures in a dark room. The enclosures were identical except for internal lights. Testing began by placing the rat in one of the two boxes. After 40–60 s, the light cues were switched, effectively ‘teleporting’ the animal to the other environment. To examine the evolution of network activity, population vectors were defined from the firing rates of all simultaneously recorded cells for every theta cycle before and after the switching of spatial cues. Each population vector was correlated with reference vectors defined from the activity of the same cells at the same spatial location in each box on separate baseline trials.

Full Methods and any associated references are available in the online version of the paper at www.nature.com/nature.

Received 25 November 2010; accepted 10 August 2011.

Published online 28 September 2011.

- Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl Acad. Sci. USA* **79**, 2554–2558 (1982).
- Amit, D. J., Gutfreund, H. & Sompolinsky, H. Storing infinite numbers of patterns in a spin-glass model of neural networks. *Phys. Rev. Lett.* **55**, 1530–1533 (1985).
- McNaughton, B. L. & Morris, R. G. M. Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends Neurosci.* **10**, 408–415 (1987).

- Treves, A. & Rolls, E. T. Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus* **2**, 189–199 (1992).
- Muller, R. U. & Kubie, J. L. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci.* **7**, 1951–1968 (1987).
- Lee, I., Yoganarasimha, D., Rao, G. & Knierim, J. J. Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. *Nature* **430**, 456–459 (2004).
- Leutgeb, J. K. *et al.* Progressive transformation of hippocampal neuronal representations in ‘morphed’ environments. *Neuron* **48**, 345–358 (2005).
- Wills, T. J., Lever, C., Cacucci, F., Burgess, N. & O’Keefe, J. Attractor dynamics in the hippocampal representation of the local environment. *Science* **308**, 873–876 (2005).
- Colgin, L. L. *et al.* Attractor-map versus autoassociation based attractor dynamics in the hippocampal network. *J. Neurophysiol.* **104**, 35–50 (2010).
- O’Keefe, J. & Nadel, L. *The Hippocampus as a Cognitive Map* (Oxford Univ. Press, 1978).
- Gothard, K. M., Skaggs, W. E., Moore, K. M. & McNaughton, B. L. Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *J. Neurosci.* **16**, 823–835 (1996).
- Samsonovich, A. & McNaughton, B. L. Path integration and cognitive mapping in a continuous attractor neural network model. *J. Neurosci.* **17**, 5900–5920 (1997).
- Derdikman, D. *et al.* Fragmentation of grid cell maps in a multicompartment environment. *Nature Neurosci.* **12**, 1325–1332 (2009).
- O’Keefe, J. & Speakman, A. Single unit activity in the rat during a spatial memory task. *Exp. Brain Res.* **68**, 1–27 (1987).
- Gothard, K. M., Skaggs, W. E. & McNaughton, B. L. Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. *J. Neurosci.* **16**, 8027–8040 (1996).
- Skaggs, W. E. & McNaughton, B. L. Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. *J. Neurosci.* **18**, 8455–8466 (1998).
- Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M.-B. & Moser, E. I. Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* **305**, 1295–1298 (2004).
- Tsodyks, M. & Sejnowski, T. Associative memory and hippocampal place cells. *Int. J. Neural Syst.* **6**, 81–86 (1995).
- Romani, S. & Tsodyks, M. Continuous attractors with morphed/correlated maps. *PLoS Comput. Biol.* **6**, e1000869 (2010).
- McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I. & Moser, M. B. Path integration and the neural basis of the ‘cognitive map’. *Nature Rev. Neurosci.* **7**, 663–678 (2006).
- 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).
- Fenton, A. A. & Muller, R. U. Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proc. Natl Acad. Sci. USA* **95**, 3182–3187 (1998).
- Olypher, A. V., Lánský, 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).
- Jackson, J. & Redish, A. D. Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. *Hippocampus* **17**, 1209–1229 (2007).
- Kelemen, E. & Fenton, A. A. Dynamic grouping of hippocampal neural activity during cognitive control of two spatial frames. *PLoS Biol.* **8**, e1000403 (2010).
- Blumenfeld, B., Preminger, S., Sagi, D. & Tsodyks, M. Dynamics of memory representations in networks with novelty-facilitated synaptic plasticity. *Neuron* **52**, 383–394 (2006).
- Tsodyks, M. V., Skaggs, W. E., Sejnowski, T. J. & McNaughton, B. L. Population dynamics and theta rhythm phase precession of hippocampal place cell firing: a spiking neuron model. *Hippocampus* **6**, 271–280 (1996).
- Akrami, A., Liu, Y., Treves, A. & Jagadeesh, B. Converging neuronal activity in inferior temporal cortex during the classification of morphed stimuli. *Cereb. Cortex* **19**, 760–776 (2009).
- Harris, K. D., Csicsvari, J., Hirase, H., Dragoi, G. & Buzsáki, G. Organization of cell assemblies in the hippocampus. *Nature* **424**, 552–556 (2003).
- Colgin, L. L. *et al.* Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* **462**, 353–357 (2009).

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

Acknowledgements We thank R. Skjerpeng, A. M. Amundsgård, K. Haugen, K. Jenssen, E. Kråkvik and H. Waade for technical assistance. The work was supported by the 7th Framework Programme of the European Commission (‘SPACEBRAIN’, grant agreement no. 200873), an Advanced Investigator Grant to M.-B.M. from the European Research Council (‘ENSEMBLE’), the Kavli Foundation, a Centre of Excellence grant from the Norwegian Research Council, and research projects MSMT CR LC554, 1M0517 and AVOZ50110509 at the Academy of Sciences of the Czech Republic.

Author Contributions K.J., A.T., M.-B.M. and E.I.M. designed the study and discussed analyses and results; K.J. built the apparatus, K.J. and E.J.H. performed experiments, K.J. performed analyses; E.I.M. wrote the paper with input from all authors.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of this article at www.nature.com/nature. Correspondence and requests for materials should be addressed to K.J. (karel.jezek@biomed.cas.cz) or E.I.M. (edvard.moser@ntnu.no).

METHODS

Subjects. Six male Long Evans rats (400–500 g at implantation) were housed individually in transparent Plexiglass cages (45 cm × 30 cm × 35 cm). The animals were kept at ~90% of their initial free-feeding body weight and maintained on a 12-h light/12-h dark schedule. All testing occurred in the dark phase. The experiments were performed in accordance with the Norwegian Animal Welfare Act and the European Convention for the Protection of Vertebrate Animals used for Experimental and Other Scientific Purposes.

Electrode preparation and surgery. Neuronal ensemble activity was recorded from ensembles of CA3 or CA1 cells in rats implanted with a ‘hyperdrive’ containing 14 independently movable tetrodes assembled in a circular bundle. Tetrodes were twisted from four 17- μ m polyimide-coated platinum-iridium wires (90% and 10%, respectively; California Fine Wire Company). Electrode tips were plated with platinum to reduce electrode impedances to 120–200 k Ω at 1 kHz.

The animals were food deprived 12 h before surgery started. In four animals, anaesthesia was induced by first placing the animal in a closed glass box filled with isoflurane vapour and then giving the animal an intraperitoneal injection of Equithesin (pentobarbital and chloral hydrate; 1.0 ml per 250 g body weight). Two animals (15272 and 15273) were anaesthetized with isoflurane (induction chamber level of 4.0% while the rats were secured in the stereotaxic apparatus, with an air flow at 1,400 ml min⁻¹; isoflurane was then gradually reduced to 1–2% during the course of the surgery). Supplementary anaesthesia was given when breathing and reflexes changed. Local anaesthetic (Xylocain) was applied on the skin before making the incision. The hyperdrive was then implanted. The tetrodes were inserted above CA3 of the right hippocampus, with the centre of the bundle at anterior–posterior 3.8 mm and medial–lateral 3.0 mm relative to bregma. Jewellers’ screws and dental cement were used to secure the hyperdrive to the skull. Two screws were connected to hyperdrive ground. All tetrodes were turned after the surgery to be sure they were in the brain.

Tetrode positions. Over the course of ~3–4 weeks, the majority of the tetrodes were lowered towards CA3 in steps of 50 μ m or less while the rat rested on a towel in a large flower pot on a pedestal. Turning was slowed down when large-amplitude theta-modulated complex-spike activity appeared in CA3 at depths of approximately 3.0 mm. The tetrode depths were tweaked to get the maximal number of simultaneously recorded CA3 cells at the start of cue switching. To maintain stable recordings, the electrodes were not moved at all on the day of recording. A few tetrodes were left in CA1; data from these tetrodes were analysed separately. Two of the tetrodes were used, respectively, to record a reference signal from the corpus callosum and an EEG signal from the stratum lacunosum-moleculare.

Behavioural training procedures. The rats were trained to collect food morsels in either of two distinct 60 cm × 60 cm enclosures with 40-cm walls located in a dark curtained environment (Supplementary Fig. 2). The boxes rested on a plexiglass plate fixed 10 cm above the floor of the room. The boxes were identical except for the arrangement of a number of internal lights. Beneath the plexiglass plate of box A there was a panel of eight light-emitting white diodes (LEDs) organized into a circle (50 cm in diameter) and placed centrally under the plexiglass plate to be visible through the floor. The box was polarized by another LED at the upper edge of one of the walls. Box B was illuminated by a 60-cm-long array of green LEDs lining 40 cm of the upper edge of the wall opposite to the directional LED in A and 20 cm of one of the adjacent walls (Supplementary Fig. 13a). The LEDs were the only light source in the room.

Training occurred in four stages. During stage 1, the boxes were located next to each other, connected by a 20 × 20 cm (width × length) passageway that allowed the rat to shuttle between the boxes in order to associate each box with a different set of path integrator coordinates⁹. The rat was permitted to travel freely between the boxes for 20 min on at least three trials. Trials were separated by 20 min intervals, during which the animal rested on a towel in a pedestal outside of the curtains. At stage 2, the corridor was removed and the animal explored the boxes individually on alternating trials (3 trials in each). At stage 3, the boxes were replaced by a single box made of the same material and equipped with both sets of lights. The box was placed on alternating trials at the two original locations. When presented at the original position of box A, the set of LEDs defining A was switched on; when in the place of box B, the respective lights of B were active instead. Again, the rat was tested 3 times in each environment on alternating occasions. Finally, at Stage 4, the box was moved to a central location between the two original box locations. The animal received alternating trials with each set of lights (2 consecutive days, each day 3 pairs of trials). During all stages, at the start of each trial, the rat was taken from the flower pot on the pedestal outside of the curtains and placed, without disorientation, into the environment with the eyes gently covered by the experimenter’s palm. Between trials, the rat rested for 20 min in the flower pot. During this period, the boxes were thoroughly washed with a wet tissue and dried.

On the test day, the rat started with a 20 min rest trial in the flower pot. The animal was then tested for 10 min in each box configuration (A and B, respectively). Then, during the third trial, after 40–60 s of baseline recording in one of the configurations (for example, A), the lights were switched instantaneously to the other configuration (for example, B). Additional ‘teleportations’ were performed every subsequent 40–60 s until 10 min had passed. A resting trial in the flower pot was recorded at the end. The electrodes were then lowered deeper into the CA3 and the experiment was repeated on the subsequent day in those cases where spikes with sufficient amplitudes from new cells could be recorded.

Running was motivated by small crumbles of cookies thrown into the box at 10–20 s intervals. Three types of cookies were used: vanilla, chocolate and unflavoured. Vanilla and unflavoured were given in configuration A, chocolate and unflavoured in configuration B. The proportion between flavoured and unflavoured cookies was approximately 1:1 during stage 1 and 1:4 during the following stages. During teleportation trials, only unflavoured crumbles were offered.

Population vector analyses. Theta waves were identified from filtered local EEG traces as described in the Supplementary Methods. The evolution of unit activity over successive theta cycles before and after cue switching was estimated by defining a population vector for each theta cycle, consisting of the number of spikes fired by each cell in the array, and comparing it to a population vector for the same cells at the corresponding location in each of the reference environments (Fig. 1a, c). Boundaries between theta cycles were determined by plotting, for all cells on the entire teleportation session, the number of spikes as a function of theta phase (Fig. 1b). The boundary between successive theta cycles was then defined by the theta phase with the lowest overall firing rate (bins of 10 deg; Fig. 1b, c). For each theta cycle, the momentary population vector was correlated with reference population vectors for the same cell sample at the same location in each of the reference environments. The reference vector (expressed in Hz rather than as a spike count, normalized by the time spent at each location) was based on activity across the entire reference trial for each of the environments on the test day. In a subset of the analyses, the momentary population vector was compared with all reference population vectors in the two boxes, that is, not only those corresponding to the animal’s current location. Correlations between momentary population vectors and reference vectors were quantified by dot products as well as Pearson product–moment correlations. Dot products were calculated for all theta cycles, including those with no activity, where the dot product is zero. Pearson correlations, which cannot be defined for theta cycles with no spikes, were calculated by excluding also theta cycles with a single unit active, that is, these correlations were based on the subset of theta cycles that included at least two active cells. In a subset of the analyses, also cycles with only two active units were excluded, and in a further control, also those with only three active units were discarded (Supplementary Fig. 7).

To determine if the similarity between momentary population vectors during cue switching and population vectors in the reference environment was larger than expected by chance, we correlated population vectors from successive theta cycles in reference recordings from A and B with mean population vectors generated from activity in the same position across the entire session in the same environment or the alternative environment. When referenced to the same environment (A × A or B × B), the distribution of Pearson correlations was centred at high correlation values with a long tail of low values (Supplementary Fig. 9, top row). When referenced to the alternative environment (A × B or B × A), the distributions were centred at slightly negative values with a long tail towards infrequent high positive values (Supplementary Fig. 9, bottom row). We then determined 5th and 95th percentile values for each distribution and used these as criteria to identify flickers to the alternative representation. Individual theta cycles were defined as A-correlated if the correlation with reference environment A was above the 95th percentile for B × A (that is, more similar to A than 95% of the theta cycles in B) and if the correlation with reference environment B was below the 5th percentile for B × B (that is, more different from the reference in B than most theta cycles in B). Conversely, theta cycles were defined as B-correlated if the correlation with reference environment B exceeded the 95th percentile for A × B and if the correlation with reference environment A was lower than the 5th percentile for A × A. In separate analyses, the momentary population vectors were compared not only with reference vectors at the same location but at all 900 bin locations in the box (Supplementary Fig. 8).

The abundance of theta cycles that correlate with both reference environments (‘mixed’ theta cycles) was determined by comparing the data with spike patterns obtained by ‘shuffling’ the activity of individual units. For each theta cycle during the 10 s preceding and succeeding the cue switching, the number of spikes produced by each unit recorded during the cue switching was drawn at random from among all theta cycles recorded in the same physical location in the same environment during the corresponding period (most of these theta cycles were correlated with the current environment, some with the alternative one). The analysis was

limited to theta cycles for which a minimum of 3 theta cycles had been recorded in that particular location in the relevant period, with 'same' location defined as the same $6\text{ cm} \times 6\text{ cm}$ spatial bin of the same recording box. The shuffled population vectors can be conceived as approximating the linear combination $\mathbf{r}_S \approx \mathbf{A} \cdot \mathbf{r}_A + \mathbf{B} \cdot \mathbf{r}_B$ where \mathbf{A} and $\mathbf{B} = 1 - \mathbf{A}$ are random binary vectors (for example, (0,1,0,0,1,0,1,1,...)) indicating whether each unit was drawn from the A or B representation of that location. If the current environment is A, a particular theta cycle in the original data can express activity close to the reference vector A in that location, or to the reference vector B (for example, during a flicker event) or to a mixture of the two, or just noise or, rather frequently, no activity at all. The abundance of mixed theta cycles was determined by the un-normalized dot product DP of the population vector in each theta cycle with the reference vectors in A and B, $DP_A = \mathbf{r} \cdot \mathbf{r}_A / C$ and $DP_B = \mathbf{r} \cdot \mathbf{r}_B / C$ (just divided by the number C of simultaneously recorded units). DP values range from 0 to about 10 (note that they are expressed in Hz, as the momentary population vector includes spike counts whereas the reference vector indicates firing rates), and very rarely beyond, but are mostly clustered close to zero because of the presence of theta cycles with no

spikes or with, for example, a single spike from a unit that emits on average less than 1 spike in that physical location in both reference vectors. We arbitrarily set $DP > 1$ as the criterion for substantial correlation, and $DP > 0.5$ for half-cycles.

The speed of instantaneous transitions from one representation (for example, A) to the other (for example, B) was estimated by cross-correlating the sequence of correlations with reference environment A and the sequence of correlations with B one time-bin later, excluding, in this analysis, units with overlapping fields in the two environments and trials with less than 20 non-overlapping units (2 out of 11; Supplementary Fig. 3). Bin widths were then varied from 44 ms to 500 ms to determine the transition time that gave the largest cross-correlation. Bins were also defined by actual theta cycles, with separations at the point of minimum spike activity in the population (Fig. 1b, c) and at other phases 10 deg apart (Supplementary Fig. 3). Cross-correlation values were corrected for overlap between A and B representations by subtracting the cross-correlation at zero time lag.

Other. Recording procedures, criteria for spike sorting, construction of rate maps, analysis of theta rhythm and histological procedures are described in Supplementary Materials and Methods.