

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/121771/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Boccaro, Charlotte N., Nardin, Michele, Stella, Federico, O'Neill, Joseph and Csicsvari, Jozsef 2019. The entorhinal cognitive map is attracted to goals. *Science* 363 (6434) , pp. 1443-1447. 10.1126/science.aav4837

Publishers page: <http://dx.doi.org/10.1126/science.aav4837>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



The Entorhinal Cognitive Map is Attracted to Goals

Charlotte N Boccara*^{†‡}, Michele Nardin[†], Federico Stella, Joseph O'Neill[#], Jozsef Csicsvari*.

IST Austria Am Campus 1, A – 3400 Klosterneuburg, Austria.

5

[†]Both these authors contributed equally to this work.

*Corresponding authors. E-mail: jozsef.csicsvari@ist.ac.at, charlotte.boccara@medisin.uio.no

[‡]Current address: University of Oslo, Sognsvannsveien 9 Domus Medica, 0372 Oslo, Norway.

10 [#]Current address: Cardiff University, 70 Park Place, Cardiff, CF10 3AT, United Kingdom.

Grid cells with their rigid hexagonal firing fields are thought to provide an invariant metric to the hippocampal cognitive map. Yet, environmental geometrical features have recently been shown to distort the grid structure. Given that the hippocampal role goes beyond space, we tested the influence of non-spatial information on the grid organization. We trained rats to daily learn three new reward locations on the cheeseboard maze, while recording from the medial entorhinal cortex and the hippocampal CA1 region. Many grid fields moved towards goal location, leading to long-lasting deformations of the entorhinal map. Therefore, distortions in the grid structure contributes to goal representation during both learning and recall, which demonstrates that grid cells participate in mnemonic coding, beyond providing a simple metric of space.

One Sentence Summary:

Goal-learning leads to a local distortion of grid cell rate maps, suggesting a complex code beyond space.

To decide upon relevant behavior, individuals rely on dynamic neural representations of their world, computed from current and past experiences. The cognitive map formed by an extended network of specialized cell types coding for defined spatial features is essential for accurate navigation (1, 2). While hippocampal place cell activity is restricted to discrete, sparse place fields in specific environments, parahippocampal grid cells present multiple firing fields, arranged in regular hexagonal arrays that densely tessellate all environments (3, 4). This led to the hypothesis that grid cells provide a universal invariant metric for spatial cognition (5). As such, they were originally considered to have a narrower role than hippocampal place cells, which code for multimodal information beyond simple spatial representations (2, 6).

New data, however, suggest a more complex grid code (7-9). Non-spatial factors modulate local field firing rates without affecting the grid structure (10) while topographically organized auditory stimulus can drive grid-like structure (11). Furthermore, geometrical environmental features can influence the rigid grid structure, thus challenging the role of grid cells to provide invariant metrics (12-14). However, grid distortions could also encode more complex behavioral information. We therefore tested the influence of behaviorally-relevant information on the entorhinal cognitive map.

We trained rats to learn daily three new hidden reward locations on a cheeseboard maze, while recording simultaneously from the medial entorhinal cortex (MEC) and the hippocampal CA1 region (fig. 1A, fig. S1). This hippocampus-dependent task (15) consisted of three phases: pre-probe, learning and post-probe, where the probes verified memory retention in the absence of food rewards (methods and fig. 1B–D). This paradigm changed daily the cognitive valence of local points in an otherwise familiar environment, leading to the accumulation of CA1 place fields

around reward locations (i.e. goal remapping, (15)). This allowed us to test how rewards can be dynamically encoded in MEC neural representation during goal learning.

The majority of grid cells (80–90%) had at least one of their firing fields significantly moving towards a goal (fig. 1E–G, binomial test, $P < 0.00001$, see methods and fig. S2–3). MEC non-
5 periodic spatial and CA1 place cells showed a similar behavior (fig. 1F–G and S3, binomial test, $P < 0.0001$). This led to the accumulation of entorhinal and hippocampal strongest firing fields at goal locations (fig. 1G). Because successful learning led to a higher number of visits of reward locations (fig. 1D), we used a Linear-Nonlinear-Poisson (LNP) spiking model to verify that the reorganization of firing fields was independent of variations in trajectories, speed or heading
10 between pre- and post-probe by comparing real data to maps generated with this method (16). We also performed an additional control by downsampling the pre- and post-probe map to match occupancy in each spatial bin (see methods and fig. 1G). Conjunctive head-direction coding did not influence field movement to goal (fig. S4).

The proportion of cell with their strongest fields within goal locations increased progressively
15 during learning in both entorhinal and CA1 spatial cells (fig. S5A–C). To test whether these changes were long-lasting, we performed the same analyses on pre-probe, this time using the previous day’s goal locations. MEC cells retained the accumulation of firing fields around previous goal locations (fig. S6A–B). Unexpectedly, CA1 changes were more transient (fig. S6C), suggesting differences in memory trace lability between those two regions, in line with models
20 arguing for faster plasticity in CA1 than in MEC (17). We also observed faster CA1 plasticity during learning by scoring the development of goal representation as a function of firing related to goal vicinity (see methods and fig. S5D–G).

Given that pre-probe grid map maintained an accumulation of firing fields at the previous day's goals, we examined the effect of the local change in cognitive valence on the grid structure itself. For three animals we added the exploration of a familiar open field, which acted as a control environment without valence bias (see methods). Grid scores were significantly higher in the open field compared to both probes (fig. 2A–B; 1-way analysis of variance (ANOVA), $P < 0.00001$), independently of differences in spatial sampling, trajectories, speed or heading (fig. 2B and methods; 1-way ANOVA, $P < 0.00001$).

To test whether the grid score drop resulted from local map distortions around goals, we used a Laplacian of Gaussian (LoG) filter to detect individual fields independently of their peak firing rate (see methods and fig. S7; (9, 10)). While there were no significant differences in the number of fields, their size, spacing or ellipticity after learning (fig. S8), we found a significant decrease in the mean distance between field center and closest reward (fig. 3A–B; Kolmogorov-Smirnov (KS) test, $P = 0.01472$) and an increase in the number of grid fields near goals (fig. 3C; Fisher's exact test, $P = 0.0145$).

To determine which parameters contributed to grid field movement, we calculated the “strength of attraction” of each field towards a goal and correlated it against spatial parameters and behavioral performance (fig. 3D–E, S9). The strongest correlation was found with the pre-probe distance to goal location: the closest fields being generally subjected to a strong attraction while fields over a distance of ca. 30 cm showing little detectable attraction (fig. 3D–E). Most attracted fields moved towards the closest goals and the most visited post-probe goal was pulling the strongest attraction (fig. S9). Different fields of a given cell could be attracted to different goals, depending on their relative position to the goals (fig. S7H–K). We observed a weak increase in the peak grid field firing rate in post-probe, however this increase was not associated with goal locations and there

was no correlation with strength of attraction (fig. S10). Finally, the increase of activity at goals correlated with memory retention (correlation analysis, $r=0.77$, $P=0.002$, fig. 3F). There was no significant difference of in the strength of attraction between CA1 and MEC fields, however pre-probe field distance influenced only MEC fields (fig. S11). To test whether the local movement of firing fields towards the closest goals could explain the grid score drop in probe sessions, we applied a movement towards imaginary goals on grid fields recorded in the open field, following the distribution of movements determined after learning. This resulted in a grid score drop was analogous to that observed between open fields and cheeseboard environments (fig. 3G; 1-way ANOVA, $P=0.0003$).

We subsequently examined the reorganization of place-related assemblies at the population level (fig. 4 and S12–13). Both MEC and CA1 cell assemblies showed a significant reduction in vector similarity between pre- and post-probe as compared to the intrinsic variability in comparing the two halves of the pre-probe (fig. S12B; t-test, all $P_s < 0.0001$). This was independent of changes in spatial sampling (fig. S12B; t-test, all $P_s < 0.0001$). MEC population vector similarity across pre/post probe sessions were weaker around goal locations than away from them (fig. S12C; t-test, $P=0.0003$) and positively correlated with distance from the goal locations (fig. 4E; $P < 0.0001$). In contrast, the reorganization of CA1 population vectors between pre- and post-probes did not exhibit significant positive correlation with goal distance (fig. 4E; $P=1$).

Given that goal remapping of individual MEC and CA1 spatial cells occurred incrementally during learning (fig. S5), we next examined how the expression of assemblies dynamically shifted toward post-probe goal representation. We computed the Fisher's z-scored correlation coefficients of the population activity in 125 millisecond time bins, with the population vectors representing the current location of the animal in the pre- and post-probes (see example session on fig. 4B and S13).

As learning progressed, similarity to post-probe representation increased both in MEC and CA1 cell assemblies. Yet, when examining the fine temporal structure of assembly expression during learning, we observed a fast-paced flickering between pre- and post-probe representations. The distribution of z-scores were then compared with a control distribution obtained by a cell-id shuffling procedure (fig. 4C). Significant differences were detected for both CA1 and MEC (KS test, $P < 0.00001$ for all combinations). Moreover, we observed a significantly heavier-tailed distribution of real z-scores compared to the control shuffled ones (binomial test, $P < 0.0001$). Real data, therefore, showed a stronger tendency to have extreme values, suggesting that there were no intermediate representations but rather flicker between the two competitive representations of pre- and post-probes. The MEC flickering we observed was reminiscent of previously reported CA1 flickering (18, 19). Given that MEC vector correlation increased with goal distance (fig. 4D), we restricted flickering analyses to goal locations. The distribution of flickering scores was significantly different and generally shifted towards more positive values (fig. S12–13, KS and Mann-Whitney U tests, $P < 0.0001$). We also observed inter-regional differences in goal flickering dynamics with CA1 reaching a plateau faster than MEC during learning (fig. S12–14).

Goal learning can lead to the local and long-lasting distortion of the entorhinal spatial maps. This demonstrates the influence of non-geometrical cognitive factors onto the grid structure itself. These findings support emerging theories proposing that the grid pattern carries a broader organizational role for both spatial and non-spatial information in more complex and naturalistic behaviors (20, 21). Grid structure distortions have been recently linked to the geometrical features of the recording environment (12–14), which may reflect distorted perception of space. Here, we found evidence for a grid code at the structural level that goes beyond simple metrics: individual grid fields moved towards newly-learned goal locations, leading to the deformation of the grid

map, independently from variations of spatial sampling, trajectories, speed or heading inherent to our behavioral paradigm.

Field attraction strength to a goal was proportional to the original goal-field distance, locally constraining the deformation of the entorhinal spatial representation. While local remapping took place in MEC, CA1 reorganized through global remapping. Moreover, field reorganization towards goal location was maintained overnight for MEC but not for CA1 although both maintained reorganized fields after learning (up to 2h). This argues for a higher lability of CA1 spatial memory traces compared to MEC cells (22). The role of the hippocampus in goal encoding was recently highlighted by reports of a subpopulation of CA1 neurons with an angular tuning for goal direction (23, 24). While our results are consistent with the role assigned to CA1 in computational models of goal directed navigation (20, 25), differences in goal coding between CA1 and MEC require updates in the current models.

Finally, we showed that assembly expression of different goal-related context rapidly alternated – flickered – in the MEC during learning, similar to CA1 flickering (18, 19). The simultaneous expression of the old and new MEC representations in the same trial periods, suggests that multiple maps can be stored in the MEC. The apparent absence of intermediate representations in MEC has implications as to how new/modified maps could be dynamically encoded and (re)organized in the CA1-MEC circuits during learning. This may open new avenues of computational research as to its role in inferential reasoning associative memory.

References and Notes:

1. J. O'Keefe, L. Nadel. *The Hippocampus as a Cognitive Map* (Oxford Univ. Press, 1978).
2. E. I. Moser, M. B. Moser, B. L. McNaughton. *Nature Neurosci.* **20**, 1448–1464 (2017).
3. T. Hafting, M. Fyhn, S. Molden, M. B. Moser, E. I. Moser. *Nature* **436**, 801–806 (2005).
- 5 4. J. O'Keefe, J. Dostrovsky. *Brain res.* **34**, 171–175 (1971).
5. E. I. Moser, M. B. Moser. *Hippocampus* **18**, 1142–1156 (2008).
6. H. Eichenbaum, P. Dudchenko, E. Wood, M. Shapiro, H. Tanila. *Neuron* **23**, 209–226 (1999).
7. E. Marozzi, L. L. Ginzberg, A. Alenda, K. J. Jeffery. *Cereb. Cortex* **25**, 4619–4627 (2015).
8. P. A. Lipton, J. A. White, H. Eichenbaum. *J. Neurosci.* **27**, 5787–5795 (2007).
- 10 9. G. W. Diehl, O. J. Hon, S. Leutgeb, J. K. Leutgeb. *Neuron* **94**, 83–92 (2017).
10. R. Ismakov, O. Barak, K. Jeffery, D. Derdikman. *Curr. Biol.* **27**, 2337–2343 (2017).
11. D. Aronov, R. Nevers, D. W. Tank. *Nature* **543**, 719–722 (2017).
12. T. Stensola, H. Stensola, M. B. Moser, E. I. Moser. *Nature* **518**, 207–212 (2015).
13. J. Krupic, M. Bauza, S. Burton, J. O'Keefe. *Science* **359**, 1143–1146 (2018).
- 15 14. C. Barry, R. Hayman, N. Burgess, K. J. Jeffery. *Nature Neurosci.* **10**, 682–684 (2007).
15. D. Dupret, J. O'Neill, B. Pleydell-Bouverie, J. Csicsvari. *Nature Neurosci.* **13**, 995–1002 (2010).
16. K. Hardcastle, N. Maheswaranathan, S. Ganguli, L. M. Giocomo. *Neuron* **94**, 375–387 (2017).
17. C. Renno-Costa, A. B. L. Tort. *J Neurosci.* **37**, 8062–8076 (2017).
18. D. Dupret, J. O'Neill, J. Csicsvari. *Neuron* **78**, 166–180 (2013).
- 20 19. K. Jezek, E. J. Henriksen, A. Treves, E. I. Moser, M. B. Moser. *Nature* **478**, 246–249 (2011).
20. D. Bush, C. Barry, D. Manson, N. Burgess. *Neuron* **87**, 507–520 (2015).
21. M. M. Garvert, R. J. Dolan, T. E. Behrens. *Elife* **6** (2017).
22. T. Hainmueller, M. Bartos. *Nature* **558**, 292–296 (2018).
23. A. Sarel, A. Finkelstein, L. Las, N. Ulanovsky. *Science* **355**, 176 (2017).
- 25 24. J. L. Gauthier, D. W. Tank. *Neuron* **99**, 179–193 (2018).

25. J. R. Hinman, H. Dannenberg, A. S. Alexander, M. E. Hasselmo. *J. Neurophysiol.* **119**, 2007–2029 (2018).
26. J. O’Neill, C. N. Boccara, F. Stella, P. Schoenenberger, J. Csicsvari. *Science* **355**, 184 (2017).
27. J. Csicsvari, H. Hirase, A. Czurko, A. Mamiya, G. Buzsaki. *J. Neurosci.* **19**, 274–287 (1999).
- 5 28. K. D. Harris, H. Hirase, X. Leinekugel, D. A. Henze, G. Buzsaki. *Neuron* **32**, 141–149 (2001).
29. C. N. Boccara *et al.* *Hippocampus* **25**, 838–857 (2015).
30. P. Quilichini, A. Sirota, G. Buzsaki. *J. Neurosci.* **30**, 11128–11142 (2010).
31. B. Dunn, D. Wennberg, Z. Huang, Y. Roudi. *bioRxiv* 101899 (2017).
32. C. N. Boccara *et al.* *Nature Neurosci.* **13**, 987–994 (2010).
- 10 33. H. Stensola *et al.* *Nature* **492**, 72–78 (2012).

Acknowledgments:


We thank Dori Derdikman for comments on an earlier version of the manuscript. **Funding:** This work was supported by a European Research Council Consolidator Grant (281511) and the European Union’s Horizon 2020 (Grant 665385). **Author contributions:** C.N.B and J.C. designed and implemented the study and wrote the manuscript. C.N.B performed the experiments. J.O’N helped with the implementation of the study. C.N.B and M.N. performed cluster cutting. C.N.B., M.N. and F.S. planned the analyses. M.N. and F.S. analyzed the data. M.N. wrote the methods. All authors discussed the results and contributed to the manuscript. **Competing interests:** The authors declare no conflicts of interest. **Data and materials availability:** Original data and programs were stored in the scientific repository of the Institute of Science and Technology Austria: <https://datarep.app.ist.ac.at/>.

Supplementary Materials:

Materials and Methods

Figures S1-S14

References (26–33)






Fig. 2 Grid score degradation on the cheeseboard.

(A) Example of MEC grid cell, which exhibits its highest grid score on the open field and a degraded score on pre- and post-probes. Top: rate maps with maximum firing rate in red and lowest in blue with peak firing rate (Hz) on the upper right corner. Bottom: corresponding spatial autocorrelogram maps, range from +1 (red) to -1 (blue), with the grid score noted on the upper right corner. (B) Average grid score (\pm SEM) across MEC cells in open field (OF: solid cyan), pre-probe (Pre: solid blue) and post-probe (Post: solid red) (1-way ANOVA, $P < 0.00001$). Dashed bars: downsampled control data (1-way ANOVA, $P < 0.00001$). Light cyan and empty bar: control data obtained with LNP spiking model in open field (LNP-OF) (1-way ANOVA, $P < 0.00001$).




Fig. 3 Stronger attraction of MEC fields closest to goal locations.

(A) Example of grid fields movement from pre-learning (blue) to post-learning (red); overlay (right). Top: animal path (grey), spikes (colored dots), goals (black dots). Bottom: field detected (colored circles) detected by LoG filter; grey-scale rate maps. (B) Density of fields as a function of distance to closest goal in pre- and post-probe; KS test, $P=0.01472$. (C) Proportion of fields at goals in pre- and post-probes; Fisher's exact test, $P=0.0145$. (D) Scatter plot of attraction strength against pre-probe goal distance. Running average (purple line, shaded SEM). Significance threshold: dashed line (t-test on each 10 cm window below 29 cm: $P<0.05$). (E) Running average of mean inferred attraction as a function of distance to closest goal. Spearman correlation analysis, $r=-0.192$, $P=0.01225$; standard-deviation (shadow). (F) Memory retention (time at goals) against normalized mean firing rate increase at goals. One dot per session. Line: regression analysis, $r=0.7$, $P=0.011$. (G) Grid score in open field (left) and distorted open field (right, generated following real-data field movement distribution); 1-way ANOVA, $P=0.0003$.




Fig. 4 Flickering in MEC and CA1 between co-existing representations.

(A) Schemas illustrating population-vector computation. (B) Example of progression of z-scored correlation coefficients of ensemble activity across paradigm with pre-probe (negative, blue) and post-probe (positive, red) population vectors in MEC (up) and CA1 (down). Yellow line: smoothed score progression; black dashed line: regression analysis. (C) Distribution of z-scores in MEC (green), CA1 (brown) and cell-ID shuffling (black); KS-test: MEC vs. shuffling: $P < 0.00001$, CA1 vs. shuffling: $P < 0.00001$. The shuffling distributions did not differ: KS-test, $P > 0.1$. (D) Running average of population vector similarity as a function of distance to goal. MEC: green; CA1: brown. 1-sided-t-test regression analysis: MEC: $r = 0.1181$, $P < 0.00001$; CA1: $r = -0.00227$, $P = 1$. Spearman correlation analysis: MEC: $r = 0.10415$, $P < 0.00001$, CA1: $r = -0.02331$, $P = 0.02369$. Shadows: standard-deviation.

5

Supplementary Materials for

The Entorhinal Cognitive Map is Attracted to Goal

10

Charlotte N Boccara, Michele Nardin, Federico Stella, Joseph O'Neill, Jozsef Csicsvari.

correspondence to: jozsef.csicsvari@ist.ac.at, charlotte.boccara@medisin.uio.no

15

This PDF file includes:

Materials and Methods
Figs. S1 to S14

20

Materials and Methods

Subject and surgeries

Four male long-Evans rats (3–5 months old, 300–400 g at implantation, housed and food deprived as described previously (26) were implanted with 16 independently movable tetrodes under deep anesthesia using isoflurane (0.5%–3%), oxygen (1–2 l/min), and an initial dose of buprenorphine (0.1 mg/kg). Our tetrodes were arranged in a drive with two bundles: one targeting the CA1 region of the hippocampus and the other the medial entorhinal cortex (MEC). All procedures involving experimental animals were carried out in accordance with Austrian animal law (Austrian federal Law for experiments with live animals) under a project license approved by the Austrian Federal Science Ministry.

Each tetrode consisted of four 12 μ m tungsten wires twisted together with their tip gold plated to reduce electrode impedance to 200–600 k Ω . Surgical implantations of electrodes were conducted following procedures previously described (26). In brief, MEC and CA1 bundles consisted of 8 tetrodes each. MEC bundles were positioned on the medio-lateral (ML) axis at [4.0 mm to 5.7 mm] lateral to the midline and, on the antero-posterior (AP) axis, [-7.4 mm to -8.8 mm] posterior to bregma. CA1 coordinates were ML [2.7 mm to 3.7 mm] and AP [-2.7 mm to -3.7 mm]. Electrodes were implanted 0.9 mm below dura at a 10-degree angle, relative to the skull.

Data acquisition

General data acquisition procedures have been described previously (26). In brief, after a recovery period of 1 week, tetrodes were progressively lower in maximum steps of 200 μ m per day until reaching CA1 pyramidal layer and MEC superficial layers while the animals were trained on the cheeseboard maze (see behavioral training procedures below).

Apparatus

The cheeseboard maze is similar to what has been described previously (15). It consisted of a PVC circular board (120 cm in diameter, 2 cm in thickness) with a total of 177 food wells (2.5 cm in diameter, 1.5 cm in depth) drilled into the surface of the maze in evenly spaced parallel rows and columns (8 cm between the centers of each well). A PVC grey start-box (27 cm long, 19 cm wide and 59 cm high) was equipped with a door (35 cm high and placed along the edge of the board perpendicular to the rows of food wells. The top of the box was open to allow tracking the animal inside. A small glass-made cup (2.5 cm in diameter) was placed inside the start-box.

Three out of four animals were also trained on a familiar open field (OF) arena very similar to the cheeseboard maze and, in most sessions, consisting of a plain circular environment of the same diameter and the same material, without food wells drilled in it. In two sessions, we used a rectangular environment. The same start box was used for both environments.

All environments were surrounded by black curtains and polarized by a 30 cm wide white cue card attached to the curtain.

Training procedures

Training procedures were similar to what has been described previously (15). In brief, rats were pre-handled before surgical procedures. Following the postoperative recovery period, they were food-deprived so that their weight was reduced and maintained at 85% of their age-matched preoperative weight. They were first habituated to retrieve pellets in the open field while electrodes were lowered to the region of interest. Once coverage of the open field was satisfactory, animals were trained to retrieve hidden pellets on the cheeseboard. First, the rat was allowed to freely

explore the whole cheeseboard for at least 30 min for 3 days. Then, the rat was trained to chase for food-rewards and come back to the start-box. Three groups of visible food pellets (MLab rodent tablet 45 mg, TestDiet) were spread out on the surface of the cheeseboard maze while the rat was inside the start box. For each trial, the door was temporarily opened, the animal was allowed to exit the box and retrieve all the rewards while another additional reward was placed in the glass-made cup situated within the start-box. Once all the rewards had been collected, the door was re-opened, and the rat was gently conducted back to the start box to find and consume the additional food reward within the start box. That procedure was repeated until the rat started to return back consistently on its own after having collected all the rewards within the board (~3–4 days). A similar procedure was applied over the following days, this time with three hidden rewards (i.e., one food pellet per baited location) within the cheeseboard maze (~2–3 days). The same baited locations were used from one day to the other.

To prevent the use of an odor-guided search strategy during these experiments, food pellet dust was scattered across the maze before each experiment, the board was periodically wiped (using the towel used to handle the rat daily) and the board was rotated relative to the start-box between learning trials and between rest and probe sessions. This initial phase of the experiment ended when the rat was familiar with the whole procedure.

Behavioral paradigm: cheeseboard spatial memory test

The animals were housed in a separate holding room and were taken to the recording room each day prior to the experiments. Each daily experiment consisted of a sequence of five recording sessions in the following order: a probe test (“pre-probe”), an immobility/sleep rest session (“pre-rest”), a learning session, an immobility/sleep rest session (“post-rest”) and a probe test (“post-probe”). The two probe tests (15~20 min) were never rewarded. After both the pre-probe and the learning sessions, rats were allowed to settle down within the start box for the rest sessions (~25 min). During the learning session, rats were given successive trials (~40 trials) to find the three hidden rewards placed in randomly selected food wells.

Three out of four animals were also exposed daily to a familiar open field (OF) arena. In that case the rat was let free to explore (15~20 min) the familiar open field environment before the beginning or after the end of the above described experimental procedure. Pellet dust and pellet crumbs were scattered on the disk.

In all, 13 sequences of probe–rest–learning–rest–probe in four animals were analyzed, with a total of 490 CA1 pyramidal cells and 262 MEC principal cells. 9 out of 13 sessions (3 rats) also had the OF exploration before the pre-probe or after the post-probe exploration, with a total of 370 CA1 pyramidal cells and 209 MEC principal cells.

DATA ANALYSIS

Spike sorting

The spike detection in the local field potential and sorting was performed as previously described (27). Action potentials were extracted by first computing power in the 800–9000 Hz range within a sliding window (12.8 ms). Action potentials with a power of >5 SD from the baseline mean were selected and spike features were then extracted by using principal components analyses. The detected action potentials were then segregated into putative multiple single units by using the Klustakwik automatic clustering software (28) (<http://klustakwik.sourceforge.net/>). These clusters were then manually refined by a graphical cluster cutting program (27). Only units with clear refractory periods in their autocorrelation and well-defined cluster boundaries were used for

further analysis. We further confirmed the quality of cluster separation by calculating the Mahalanobis distance (28) between each pair of clusters. Periods of waking spatial exploration, immobility, and sleep were clustered together and the stability of the isolated clusters was examined by visual inspection of the extracted features of the clusters over time. Pyramidal cells and interneurons in the CA1 region were discriminated by their autocorrelations, firing rate (average 0.1~5 Hz) and waveforms, as previously described. Putative MEC principal cells were identified by their firing rate: they had to keep a stable average firing over the entire recording day between 0.1 and 7 Hz. In this way we were able to identify the activity of 262 putative excitatory MEC neurons (209 with OF), as well as 490 CA1 pyramidal units (370 with OF).

Histology and reconstruction of recording positions

Electrodes were not moved after the final recording session. The rats were killed with an overdose of pentobarbital and were transcardially perfused with 0.9% saline (wt/vol) followed by 4% formaldehyde (wt/vol). The brains were extracted and stored in 4% formaldehyde. At least 24 h later, the brains were quickly frozen, cut in sagittal sections (30 μ m) using a cryostat, mounted and stained with cresyl violet (Nissl). Every section in the area of the tetrode trace was retained. The positions of the tips of the recording electrodes were determined from digital pictures of the brain sections. The laminar locations of the recording electrodes in MEC and CA1 were determined on the basis of cytoarchitectonic criteria (29), as well local field potential sharp-wave response during rest (30).

Behavioral performance

Behavioral performance was calculated offline using the animal's position records from the tracking data. Learning performance was assessed by constructing learning curves based on the time and the distance travelled to retrieve all three rewards for each trial. Since the baseline time and distance changed daily, we normalized the learning curves by each day's baseline performance (last 20 trials) and then averaged them pointwise in order to compute an overall mean learning curve. Memory retention performance was assessed during the first 5 min of each probe by scoring the time spent in the goal areas, defined as 15 cm in diameter centered on the learned bait locations. In subsequent analysis, the learning was subdivided in learning blocks, each consisting of seven trials and an overlap of two with the neighboring one (fig. S5H).

Occupancy maps

Position estimates were based on tracking the middle positions between LEDs on the head stage. The x-y plane of the cheeseboard was divided into bins of 3 cm \times 3 cm and occupancy-maps were calculated during exploratory epochs (speed >3 cm/s) measuring the amount of time spent in each spatial bin by the animal, based on the tracking data. The number of periods ($\tau = 25.6$ ms) spent in each bin were counted using a Triweight Kernel (31) with a bandwidth of $\sigma = 3$ cms and centered at the center of each spatial bin x_b :

$$occ(x_b) = \sum \tau K(x_b \vee x)$$

where

$$K(x \vee \xi) = \frac{4}{9\pi\sigma^2} \left[1 - \frac{\|x - \xi\|^2}{9\sigma^2} \right]^2$$

Spatial firing rate maps

The x-y plane of the cheeseboard was divided into bins of 3 cm × 3 cm and rate-maps were calculated during exploratory epochs (speed >3 cm/s) by dividing the number of spikes recorded in each bin by the occupancy and then smoothed with a Gaussian filter with a standard deviation (SD) of two bins. Bins with less than 250 ms occupancy time were not considered. The number of spikes in each bin were counted using a Triweight Kernel with a bandwidth of 3 cm. The peak rate was defined as the rate in the bin with the highest rate in the firing rate map. Normalized firing rate maps were the original firing rate maps divided by the peak firing rate (unless otherwise specified). Z-score normalized maps were the original firing rate maps to which the mean firing rate was subtracted and divided by the map SD.

Sparsity measure and spatially selective cells labelling

Sparsity measure represents the proportion of the environment in which a cell fires, corrected for occupancy time. It is formally defined as

$$\frac{(\sum P_i R_i)^2}{\sum P_i R_i^2}$$

where P_i is the probability of the rat occupying bin i and R_i is the firing rate in bin i . Hippocampal place cells were screened for their spatial tuning using a sparsity value of no more than 0.3, whereas entorhinal spatially modulated cells were selected using a sparsity threshold of 0.6. This discrepancy is due to the fact that entorhinal cells usually present multi-field firing patterns, giving rise to a higher sparsity score even when spatially modulated.

Shuffling procedure for cell labelling

To keep the firing structure of each cell while disrupting the link between firing and spatial position we used a spike list wrapping shuffling procedure (32). In brief, considering a spike train with time span $[0, \text{END}]$, a random time between 20 sec and $\text{END} - 20$ sec was picked and added to all the spiking times. The spiking times which resulted above END were brought to the beginning by subtracting END , and the new ‘wrapped’ spike train was used to compute a new rate map. This was repeated 200 times for each cell in each environment independently.

Spatial autocorrelograms

The autocorrelogram represents the map of shifted self-coherence and was calculated as in (32). For any regularly spaced grid of gaussian bumps one expects to obtain again the same regular grid as autocorrelogram, because the coherence will be high when the shift lets the peaks overlap again.

Grid score

The grid score represents the amount of rotational coherence of a certain map. It was calculated similarly to (33): From a spatial rate autocorrelogram whose center was excluded, we considered the Pearson correlation of the autocorrelogram rotated by 30, 60, 90, 120 and 150 degrees (± 3 degrees offsets). Only bins closer to the center than a radius s were considered. The grid score using this particular radius s , was defined as the difference between the average of the maximum correlations around 60 and 120 degrees (± 3 degrees offsets) and the average of the minimum correlations around 30, 90 and 150 degrees (± 3 degrees offsets). Eventually, the final grid score of the cell was defined as the maximum grid score over values of s ranging from twenty to forty bins, computed at intervals of one bin width. A small amount of ellipticity (among the axes ratios 1, 1.1 and 1.2) was also allowed in order to correct for possible deformations in the grid structure (33).

Grid cells labelling

A cell was labelled as grid cell only when its grid score was higher than the 95th percentile of the distribution of scores coming from a spike list wrapping-shuffling procedure (see above). The threshold was calculated separately for each cell in each single environment. The cells were classified as grid cells based on their score in the open field familiar and stable environment (3 out of 4 animals). The scores in pre-probe were used only when the open field data was not present (1 out of 4 animals). With this procedure a total of 56 cells were identified.

Head-direction modulation

To determine whether cells had their activity modulated by head-direction, we computed the mean vector length of the angular spiking distribution (32). A cell was labelled as directional only when its mean vector length was higher than the 95th percentile of the distribution of scores coming from a spike list wrapping-shuffling procedure (see above). The threshold was calculated separately for each cell in each single environment.

Accumulation of firing at goal locations

This analysis tested for high spiking activity close to the goal locations. The bins with high firing rate were detected as bins with an intensity above 80% of the peak. The cells with high firing within 15 cm from any goal location were counted in pre-probe rate maps and post-probe rate maps. Maps were computed using the activity of the first 10 minutes of each probe, in order to standardize the time considered across animals and sessions. The increase between pre-probe and post-probe was tested using a Fisher's exact test on contingency tables. The same procedure was used during learning on rate maps computed during trial blocks (as described above, 7 trials with an overlap of 2) and averaged across cells.

Downsampling procedure

To verify that the observed effects were not due to the fact that the rat spent more time close to goal locations in post-probe, we used a random downsampling procedure.

Each cell was associated with a list (*position_t*, *spike_t*), where *position_t* represents the (x,y) coordinate of the animal at a particular time point and *spike_t* the number of spikes the cell emitted in that particular time bin. Each point was then associated to its own spatial bin in a map with 10 x 10 spatial bins (15 cm size).

To correct for uneven sampling between pre and post-probe, we randomly sampled while allowing repetitions, for each map, a number of events corresponding to the minimum occupancy of the two. Using the downsampled data, we calculated new rate maps using the same procedures as described above. This was repeated 200 times for each cell.

Linear–nonlinear model and Poisson simulations

To assure that the change observed in the rate maps were not due to behavioral differences, we used a linear–nonlinear (LN) model for the entorhinal cells that takes into consideration position, heading and speed of the animal. The model has been used for entorhinal cells by (16) and is described there in detail. In brief, this model has been used to quantify the dependence of spiking on a combination of variables (position, heading and speed): it allows to estimate the firing rate of a cell during a time bin as an exponential function of the sum of the relevant value of each variable projected onto a corresponding set of parameters. Once the LN model was fitted to the data observed in pre-probe, the expected firing rate were computed based on the behavior of post-probe. Using the instantaneous expected firing rate, we simulated the firing of a cell in post-probe as a

non-homogeneous Poisson process (31). This procedure was repeated 200 times for each cell and is referred through the text as linear-nonlinear Poisson spiking (LNP) model.

Firing-by-vicinity score

The firing-by-vicinity score correspond to the sum of the firing field rates of cells near goals during learning. To obtain this score we first constructed a two-dimensional kernel function by summing three 2D Gaussian kernels with a SD of 30 cm and centered at each goal. We then multiplied the normalized firing rate map of each cell with the two-dimensional kernel function. Finally, we summed up the kernel-multiplied spatial firing rates. Thus, we calculated, for each cell, their firing rates near the goals, taking into consideration those rates in the immediate vicinity and less those further away. During learning, we computed the score for each cell on a sliding window of seven trials in order to have a better temporal resolution and compared it across cell types.

Firing field detection

It has been recently shown that firing fields of grid cells present different firing intensities, so a thresholding procedure is not suitable to detect grid fields (31). To detect grid fields independently of their firing intensities, we employed a Laplacian-of-Gaussian (LoG) filter blob detection algorithm (31). This algorithm consists of convolving the image with a LoG kernel. The Laplacian filter measures the local curvature of a surface and is defined as the sum of the unmixed second derivatives of a (possibly multidimensional) real function and can be thought in 2D as a measure of local curvature of a surface. The convolved image emphasizes areas of high curvature of a gaussian-smoothed version of the image and will have negative peaks where the peaks of the fields are (fig. S7). The gaussian kernel, used before the Laplacian filter, had the standard deviation of a gaussian bell fit on the central peak of the autocorrelogram. This measure was also used as the average dimension of the fields of a map. The detected peaks were then filtered using the following criteria:

- Intensity: the intensity had to be higher than the 75th percentile of the distribution of intensities of the map.
- Overlap: if two fields overlapped, only the one with the strongest intensity was considered.
- Distance to the border: in order to avoid “half fields” or distorted fields close to the borders, all the fields detected less than 10 cm from the border were discarded.

Firing fields matching

To study the movement of individual grid fields between pre and post-probe, we paired them based on their spatial overlap and distance between centers. Once the fields were detected in pre- and post-probe, the pairing worked as follows:

- We calculated the distance of each field in pre-probe to each field in post-probe.
- Fields were matched starting from the pair with the biggest overlap (smallest distance). Once matched they were then excluded.
- The fields were paired up until reaching a maximum distance of 20 cm.

The threshold of 20 cm is given by the fact that the average spacing between fields in our dataset was 35 cm and the average field size was about 15 cm.

Significant movement towards a goal

We used two methods to determine the significance of movement towards a goal of each field against LNP data with method 1 or downsampled data with method 2 (fig. 1G).

Method 1: We first simulated 100 times the pre-probe firing fields with the LNP spiking model using the real post-probe behavior. We then evaluated the number of times a firing field moved more in the LNP-generated maps than in the real maps. Next, we compared each field movement against chance on a cell-by-cell basis using a 95% significance level. Finally, we compensated for multiple comparison with a Holm-Bonferroni correction.

Method 2: This method is very similar to method 1. We randomly downsampled 100 times the pre- and post-probe rate maps in order for them to reach the same occupancy. We then evaluated the number of times a firing field moved more in the downsampled maps than in the original maps. Finally, we compared each field movement against chance on a cell-by-cell basis using a 95% significance level with a Holm-Bonferroni correction.

After we assessed whether a field moved significantly or not, we checked whether it moved towards or further away from goals.

Attraction strength

In order to quantify the movement of the fields towards the goal locations, we measured the relative movement of each field and weighted it by a directionality score.

If we denote with d_{pre} the distance of the field from a goal in pre-probe and d_{post} the same in post probe, the attraction strength was defined as:

$$att = (d_{pre} - d_{post}) / (d_{pre} + d_{post}) * | \cos(\theta) |$$

where θ is the angle between the vector of movement and the vector pointing from the field to the goal location.

Inferred attraction strength

To verify the hypothesis that attraction is stronger on fields located close to goals, we used an algorithm that tests for the best fit between pre- and post-probe when moving the firing fields towards a goal location. For each pre-probe map, we detected the firing fields and then moved them towards one (or more) goal location(s). Each field was free to move towards any goal (or combination of goals) by any percentage of the distance that separated it from the goal it was moving towards. We then measured the correlation between the post probe map and the modified pre-probe maps, using any combination of movements towards any of the wells, and selected the one that correlated best. The inferred attraction for each field was the percentage of movement of the map that fitted best the post-probe map, and we scatter-plotted it against the distance of the fields in pre-probe in order to carry out a correlation test.

Field deformation, size and spacing

To check whether the fields were more deformed on the cheeseboard, we fitted a (diagonal) 2D gaussian bell on each field detected in the open field, in pre- and post-probe; and computed a deformation score. The score was computed as the absolute value of the difference of the two standard deviations divided by the sum of the two. As such we had a score ranging from 0 to 1, where 0 means perfectly round and 1 means perfectly flat. The mean size of the fields was computed as the standard deviation of a perfectly symmetric 2D gaussian bell fitted on the central blob of the autocorrelogram, whereas the mean spacing between fields (i.e. grid scale) as the

distance between the central blob and the one outside the center with the highest intensity of the autocorrelogram.

Grid score degradation

To test whether the movement of fields observed in probes could explain the drop of grid score, we took the firing rate maps of the grid cells recorded in the open field and applied a movement on each field following the real distribution of movements from pre- to post-probe. We repeated this procedure 100 times for each cell and then compared the grid scores of the distorted maps with the original ones.

Rate remapping

To compare the extent of rate remapping between pre- and post-probes, we employed a rate remapping score computed as:

$$\text{remap_score} = \text{rate_post} - \text{rate_pre} / \text{rate_post} + \text{rate_pre}$$

We compared both the peak firing rate of each cell and also the firing rate of each detected firing field. We then compared the distribution of scores with downsampled data, as well as with the maps generated using the LNP spiking model in order to simulate the spiking using the rate modulation fitted on pre-probe and the real behavior of post-probe.

Correlation with behavioral performance

We computed the average increase of firing at goals between pre- and post-probe using the firing-by-vicinity method described above. We then correlated it with the time the animal spent around the most visited reward area in post-probe during the first 5 or 10 minutes.

Population vector

The population vector represents the simultaneous average activity of all the selected cells in a particular spatial bin, computed in a certain experimental session. In brief, all the spatial rate maps of the cells considered (here: CA1 place cells or MEC spatially selective cells) in a given session were stacked along the z-axis. Thus, the population vector in a given spatial bin is the ‘vertical’ z vector for that particular (x,y) set of coordinates (18).

Population vector similarity against distance from goals

To analyze the dependence of population vector similarity on the distance from the goals, the spatial rate maps were z-scored and then stacked along the z-axis for pre and post-probe as described above. For each spatial bin (with an occupancy of at least 250 ms in both pre- and post-probe), the Pearson correlation between the two firing rate vectors was calculated and then scatter plotted against the distance from the goal location. We then tested whether the firing rate similarity was positively correlated with the distance from the closest goal using both a one-sided (positive) correlation analysis, as well as with a Spearman correlation analysis.

Flickering

To test the effectiveness of learning from a population point of view, we compared the instantaneous firing during learning with the average activity that we observed in pre- and post-probe. The activity of each cell of the selected ensemble (CA1 pyramidal or MEC principal cells) was binned into 125 ms windows and smoothed with a 1D gaussian filter with 250 ms SD. Each time-window overlapped 100 ms with the previous. We then measured the correlation between the instantaneous population activity in each time-window during learning with the population vector activity of pre- and post-probes. In this way, we obtained two time series expressing the similarity

of the instantaneous population activity with the representation observed during pre- and post-probes. We employed a Fisher's z-transform of the correlation coefficients and took their difference. We next compared the distribution of the z-scores against the scores coming from a cell-ID shuffling procedure, and then verified whether the two distributions were significantly different using a Kolmogorov-Smirnov test. Finally, we used a binomial test to check whether the number of scores outside the 95th percentile threshold (1.645) was higher than chance. This procedure was executed only on sessions that had at least 10 CA1 and 10 MEC units (10 out of 13 sessions).

Average flickering during learning trials

To study the temporal behavior of flickering during learning, we averaged the difference of Fisher z-scored correlation coefficients during each learning trial, pooling together all the sessions.





Fig. S1. Histological verification of electrode positioning. Nissl-stained sagittal sections showing the medial entorhinal cortex and the CA1 region of the hippocampus with electrode tracks (red arrows) terminating respectively in the superficial layers (II/III) and in the pyramidal cell layer. Inserts show the entire region and the location of the delineation of the magnified regions (red boxes). (A) animal 1, (B) animal 2, (C) animal 3, (D) animal 4. The vast majority of the

entorhinal cells were recorded in the superficial layers (II/III). A minority of cells were recorded in LV (animal 3, one tetrode).






Fig. S2. Additional examples of grid cell spatial maps during pre-probe, learning and post-probe. Spatial maps of five example MEC cells showing the accumulation of activity around goal location after learning. For each cell there are three panels of maps from top to bottom. Top panels show the firing rate maps, scale on the left of the row and peak firing rate (Hz) in the upper right corner of each map. Middle panels represent the path of the animal (grey) with colored dots (blue: pre-probe; orange; learning; red: post-probe) representing the spikes of the cells. Bottom panels show the spatial autocorrelograms, scale is on the left of the row going from -1 to +1. Purple areas on the right are divided in two panels: the upper panels show the differential heat rate maps with scale bar on the right and pre-(blue); the lower panels show the spike overlays of pre-probe (blue) and post-probe (red). The red circles highlight goals pulling strong attraction.




Fig. S3. Distribution of movement relative to goal among all spatial cells recorded.

(A) Distribution of grid cells according to the significant movement of their grid fields relative to goals. Significance calculated against downsampled data (see methods). Left “only closer”: all detected moving fields significantly moved closer to goals (69%). Middle left “mixed”: detected fields presented a mixed behavior with some moving towards the goal and some not (20%), thus 89 % grid cell show goal-attraction of at least one field (fig. 1). Middle right “only away”: all detected fields significantly moved away from goals (7%). Right “no mov”: none of the detected fields moved significantly (4%). (B) Same as (A) but for entorhinal spatial cells. In the same order as above: 60%, 24%, 6%, 10%. (C) Same as (A) but for hippocampal place cells. In the same order as above: 64%, 15%, 10%, 11%. (D) Same as (A) but with significance calculated against LNP data. In the same order as above: 66%, 20%, 5%, 5%. (E) Same as (B) but with significance calculated against LNP data. In the same order as above: 68%, 18%, 3%, 11%. (F) Same as (C) but with significance calculated against LNP data. In the same order as above: 63%, 15%, 10%, 12%.




Fig. S4. Influence of head-direction (HD) modulation on field movement. (A) Conjunctive grid-by-head-direction cell proportions at goals pre-learning (blue) and post-learning (red). Fisher's exact test: $P=0.3$, $N=10$. Binomial test: $P=0.0127$, $N=10$. As in fig. 1. Note that the non-significance for the Fisher's exact test is due to the very small number of conjunctive cells. (B) As in (A) for non-directional 'pure' grid cells. Fisher's exact test: $P<0.001$, $N=46$. Binomial test: $P<0.0001$. (C) Proportions of conjunctive grid-by-head-direction cells (as in fig. S3). Left "only closer": all detected fields significantly move closer to goals (70%). Middle left "mixed": detected fields present a mixed behavior with some moving towards the goal and some not (10%). Middle right "only away": all detected fields significantly move away from goals (20%). Right "no mov": none of the detected fields move significantly (10%). Significance calculated against downsampled data. (D) As in (C) for non-directional 'pure' grid cells. In the same order: 67%, 17%, 4%, 2%.




Fig. S5. Comparative dynamics of goal-encoding emergence during learning. (A) Proportions of grid cells with strongest firing at goals across learning-blocks. Spearman correlation, $r=0.928$, $P=0.0025$. (B) Same as (A) for entorhinal spatial cells, $r=0.835$, $P=0.0193$. (C) Same as (A) for hippocampal place cells, $r=0.53$, $P=0.2193$, first 4 blocks: $r=0.99$, $P<0.00001$. Orange shadows: standard-errors. The analyses used in fig. S1G (right) and S5A–C are based on the selection of the strongest firing fields (top 20% strongest firing bins) moving from outside to inside goal locations. Yet, it did not include the firing fields with low firing rate or movement occurring outside or within goal locations. Thus, we computed a firing-by-vicinity score where we summed the firing field rates of cells near goals by multiplying the firing rate with three 2D kernel functions centered at the goal locations, see fig S1G (left) and S5D–G. (D) Evolution of firing-by-vicinity score in grid cells across behavioral paradigm. Left: average in pre-probe (blue bar \pm SEM). Middle: scores per sliding windows of seven trials (orange). Right: average in post-probe (red bar \pm SEM). Dashed line: level measured at the end of learning. Mann-Whitney U test, pre- vs. post-probe: $P<0.001$. Regression test, trials 1–15: $P>0.98$, trials 16–30: $P<0.0001$. (E) Same as (D) for entorhinal spatial cells. Mann-Whitney U test, pre- vs. post-probe: $P<0.001$. Regression test, trials 1–15: $P>0.9$, trials 16–30: $P<0.001$. (F) Same as (D) for hippocampal place cells. Mann-Whitney U test, pre- vs. post-probe: $P<0.001$. Regression test, trials 1–15: $P<0.0001$, trials 16–30: $P>0.81$. (G) Comparison of firing-by-vicinity score (after baseline subtraction) across learning (seven trials sliding windows): CA1 place cells (brown), MEC spatial cells (dark green), MEC grid cells

(light green). (H) Average distribution of time (top) and area (bottom) per learning block (as presented in fig. S5A–C).




Fig. S6. Longer-lasting map deformation in the MEC compared to CA1. (A) Proportion of grid cells around goal locations from previous day before learning (“yesterday’s goal”, light blue). To allow easy comparison, proportion of grid cells around today’s goal are shown in bright blue (before learning) and red (after learning). Fisher’s exact tests: pre-probe, old vs. new wells: $P=0.01756$, $N=43$. Dashed bars show control data obtained with downsampled data following the same color code: light blue (before learning, old goal locations); bright blue (before learning, new goal locations); bright red (after learning, new goal locations). Fisher’s exact tests: pre-probe, old vs. new wells: $P<0.00001$, $N=4300$. (B) Same as (A) but for entorhinal spatial cells. Fisher’s exact tests: pre-probe, old vs. new wells: $P=0.0047$, $N=122$; pre-probe, old vs. new wells downsampled: $P<0.00001$, $N=12200$. (C) Same as (A) but for hippocampal place cells. Fisher’s exact tests: pre-probe, old vs. new wells: $P=0.432$, $N=183$; pre-probe, old vs. new wells downsampled: $P=0.07823$, $N=18300$.




Fig. S7. Laplacian of Gaussian (LoG) based field detection and field-to-goal movement. A–D: Schematic of the field detection algorithm. (**A**) Example grid cell with uneven field rate. Up: rate map with peak rate in black. Bottom: spatial autocorrelogram with peak correlation in red. (**B**) LoG kernel to convolve the image. (**C**) Up: same rate map as in (**A**) convolved with the LoG kernel, where σ has been estimated as the standard deviation of a symmetric Gaussian bell fitted on the central peak of the autocorrelogram. Red dots: detected local minima (minimum distance of 5 bins). Bottom: Corresponding autocorrelogram of the convolved rate map. (**D**) Up: convolved rate map with the filtered field centers (fields too close to each other, with too low intensity or too close to the border were excluded). Bottom: original rate map with circles on top of the detected fields. **E–G**: Proportions of spatial fields moving towards goals (left), not moving (middle) or away from goals (right) for grid cells (**E**), spatial MEC cells (**F**) and place cells (**G**). Left bars stratification represents proportions moving towards closest (1st), second closest (2nd) or furthest (3rd) goal. **H–J**: Proportion of cells with fields presenting none (0 goal, left) or significant movement towards one (middle left), two (middle right) or three goals (right) between pre- and post-learning for grid (**H**), spatial MEC (**I**) and place cells (**J**). Note that for this analysis, one field may be attracted to several goals. **K–M**: Proportion of cells whose individual fields showed strongest attraction towards one, two or three goals for grid cells (**K**) and spatial MEC cells (**L**).

Note that for this analysis, each individual field could only be attracted to one goal. Given that the vast majority of place cells presented only one field, this analysis was not performed for them.




Fig. S8. Comparison grid properties across open field, pre- and post-probe. (A) Mean field radius across environments; one-way ANOVA, $P=0.38638$. (B) Mean grid spacing across environments; one-way ANOVA, $P=0.12$. (C) Mean field distortion-score across environments; one-way ANOVA, $P=0.31242429$. (D) Mean number of fields per grid across environments; one-way ANOVA, $P=0.59$.




Fig. S9 Attraction score and influence of various factors on attraction pulled by each reward well. (A) Schematic illustrating attraction score computation (see methods for details). Grid fields mean attraction strength to wells ordered by: (B) attraction strength, (C) time spent at the well during post-probe, (D) distance of the well to the start box, (E) time to reach a well during learning, (F) place in the sequence of learning and (G) time spent at the well during learning. First ordered well is dark purple (right bar), second well medium purple (middle bar) and third well light purple (left bar). Significance given by one-way ANOVA, B: $P < 0.00001$ (***), $P = 0.0363$. C: $P = 0.036$ (*), D: $P = 0.225$ (NS: Non-significant), E: $P = 0.301$ (NS), F: $P = 0.124$ (NS), G: $P = 0.119$ (NS).




Fig. S10 Rate remapping absence of influence on goal remapping. Rate remapping scores were computed for each firing field independently as the difference of firing rate in post- and pre-probe divided by the sum. (A) Rate remapping scores between pre- and post-probe in real data (purple bars) or in LNP post-probe maps (blank bars) generated using post-probe behavior and LNP spiking model based on pre-probe activity. Note a significant positive tail showing rate remapping for real data through an increase of rates in post-probe. Rank-sum test: $P < 0.0001$, Levene test for Variance: $P < 0.0001$. (B) Same as in (A) for [pre- vs. post-probe] real data (purple bars) compared to control [pre- vs. post-probe] downsampled data (blank bars). Note the absence of significant differences between the two distributions showing the independence of rate remapping from spatial sampling. Rank-sum test: $P = 0.87$, Levene test: $P = 0.39$, KS test: $P = 0.76$. (C) Same as in (A) for [pre- vs. post-probe] real data (purple bars) compared to intrinsic variability in first vs. second half of pre-probe (blank bars). Note a significant positive tail showing positive rate remapping from pre- to post-probes, larger than the intrinsic rate variability within a session. Rank-sum test: $P < 0.023$, Levene test: $P = 0.0021$. (D-H) Regression analyses showing the absence of correlation between rate remapping and various movement measures on a field-by-field basis. Regression analysis, all $P_s > 0.1$. (D) Mean attraction score as a function of rate remapping score. (E) Maximum attraction score as a function of rate remapping score. (F) Field movement as a function of rate remapping score over bins of 3 cm each. (G) Distance from well with strongest attraction as a function of rate remapping score. (H) Distance from well with strongest attraction

as a function of difference in peak firing rate. **(I)** Average coefficient of variation of firing field intensities across cells. one-way ANOVA, $P>0.2$.




Fig. S11 CA1 goal attraction strength compared to MEC: absence of influence of goal distance and weaker correlation with memory retention. **(A)** Average strength of attraction calculated for fields within a diameter of 30 cm around goals for grid (light green), MEC spatial (dark green) and CA1 place (brown) fields. one-way ANOVA, $P=0.3873$. **(B)** Same as (A) but for average strongest attraction. one-way ANOVA, $P=0.4357$. **(C)** Scatter plot of CA1 attraction strength against pre-probe's distance to goal. Running average (purple line, shaded SEM). No significant difference found from zero (t-test on each 10 cm window: all $P_s>0.05$). **(D)** Running average of mean CA1 inferred attraction as a function of distance to closest goal was not significantly different. Spearman correlation analysis, $r=0.0207$, $P=0.748$; shadow: standard deviation. **(E)** Memory retention (time at goals) against normalized mean CA1 firing rate increase at goals during the first 5 mins in post-probe. One dot per session. Red line: regression analysis: $r=0.66$, $P=0.015$. **(F)** Same as (D) except that the normalized mean firing rate increase was calculated over 10 mins. Regression analysis: $r=0.38$, $P=0.201$. Note that the correlation presented on figure 3F for MEC showed the mean firing rate increase at goals during the first 10 minutes in post-probe. Entorhinal increase of activity at goals also significantly correlated with the 5 minutes memory retention test (data not shown).




Fig. S12 Population vector and flickering analyses: focus on goal areas. (A–B) Schemas illustrating population-vector computation for grid cells (A) and place cells (B). (C) Population vector similarity in MEC. Left: first vs. second half of pre-probe (light green bar). Middle: pre-probe vs. post-probe (green bar). Right: downsampled pre-probe vs. post-probe (blank bar); t-tests: $P < 0.00001$ for all. (D) Population vector similarity in CA1. Left: first vs. second half of pre-probe (light brown bar). Middle: pre-probe vs. post-probe (brown bar). Right: downsampled pre-probe vs. post-probe (blank bar); t-tests: $P < 0.00001$ for all. (E) Local population vector similarity in MEC restricted around or outside goal locations. Pre-probe vs. post-probe for grid cells inside (left yellow-green bars) and outside (right dark green bars) goal area; t-tests: $P = 0.00033$. (F) Same as (E) for CA1 populations. Pre-probe vs. post-probe for place cells inside (left yellow bars) and outside (right dark brown bars) goal area; t-tests: $P = 0.2001$. (G) Comparison of distributions of MEC flickering z-scores in the entire environment and restricted to goal areas. KS test: $P < 0.0001$. Rank-sum test: $P < 0.0001$. (H) Same as (G) for CA1 populations. KS test: $P < 0.0001$. Rank-sum test: $P < 0.0001$. Note an increase flickering in goal areas. (I) Average of the z-scores values during learning per trial across sessions in MEC (green) and CA1 (brown). Shades represent SEM. Note scales are different for MEC and CA1. Regression test, trials 1-16: MEC: $r = 0.91$, $P < 0.0001$, CA1: $r = 0.98$, $P < 0.00001$; trials 15-30: MEC: $r = 0.93$, $P < 0.00001$, CA1: $r = 0.19$, $P = 0.46$. (J) Same as in (I) but restricted around goal locations. Regression test, trials 1-16: MEC: $r = 0.79$, $P = 0.002$, CA1: $r = 0.95$, $P < 0.00001$; trials 15-30: MEC: $r = 0.97$, $P < 0.00001$, CA1: $r = 0.09$, $P = 0.76$. (K) Same data as in (I) and (J) but scaled to allow different comparison between CA1 and MEC. Note that: though the passage from positive to negative flickering may be informative, we believe one should focus on the latency to reach a plateau.




Fig. S13. Additional examples of z-scores regression slopes. Examples of progression of z-score difference of MEC (up) and CA1 (down) populations activity similarity to pre- and post-probe representations during two example days spanning (from left to right): pre-probe, learning and post-probe. Z-score difference more similar to pre-probe are represented by blue bars and to post-probe by red bars. Yellow line: smoothed score progression; black dashed line: regression analysis.




Fig. S14. Example of z-scores regression slopes at goal location during learning. Example from fig. 4. Top: analyses performed on whole cheeseboard. Bottom: analyses restricted to goal locations only.