מכון ויצמן למדע
WEIZMANN INSTITUTE OF SCIENCE

Multiple Maps of the Same Spatial Context Can Stably Coexist in the Mouse Hippocampus

Document Version:
Accepted author manuscript (peer-reviewed)

Citation for published version:
Sheintuch, L, Geva, N, Baumer, H, Rechavi, Y, Rubin, A & Ziv, Y 2020, 'Multiple Maps of the Same Spatial Context Can Stably Coexist in the Mouse Hippocampus', *Current Biology*, vol. 30, no. 8, pp. 1467-1476. https://doi.org/10.1016/j.cub.2020.02.018

Total number of authors:
6

Digital Object Identifier (DOI):
10.1016/j.cub.2020.02.018

Published In:
Current Biology

License:
Other

General rights
@ 2020 This manuscript version is made available under the above license via The Weizmann Institute of Science Open Access Collection is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognize and abide by the legal requirements associated with these rights.

How does open access to this work benefit you?
Let us know @ library@weizmann.ac.il

Take down policy
The Weizmann Institute of Science has made every reasonable effort to ensure that Weizmann Institute of Science content complies with copyright restrictions. If you believe that the public display of this file breaches copyright please contact library@weizmann.ac.il providing details, and we will remove access to the work immediately and investigate your claim.

(article begins on next page)

Multiple maps of the same spatial context can stably coexist in the mouse hippocampus

Liron Sheintuch[1], Nitzan Geva[1], Hadas Baumer, Yoav Rechavi, Alon Rubin[2], Yaniv Ziv[2,3]

Department of Neurobiology, Weizmann Institute of Science, Rehovot 76100, Israel

[1, 2] These authors contributed equally to this work.

[3]Lead Contact

Correspondence to: yaniv.ziv@weizmann.ac.il and alon.rubin@weizmann.ac.il

SUMMARY

Hippocampal place cells selectively fire when an animal traverses a particular location, and are considered a neural substrate of spatial memory. Place cells were shown to change their activity patterns (remap) across different spatial contexts, but to maintain their spatial tuning in a fixed familiar context. Here, we show that mouse hippocampal neurons can globally remap, forming multiple distinct representations (maps) of the same familiar environment, without any apparent changes in sensory input or behavior. Alternations between maps occurred only across separate visits to the environment, implying switching between distinct stable attractors in the hippocampal network. Importantly, the different maps were spatially informative and persistent over weeks, demonstrating they can be reliably stored and retrieved from long-term memory. Taken together, our results suggest that a memory of a given spatial context could be associated with multiple distinct neuronal representations, rather than just one.

INTRODUCTION

A prevailing notion in memory research is that a given memory is supported by a single neuronal representation. In the hippocampus, such a representation, often considered the neuronal basis for a 'cognitive map', emerges from the activity of a population of neurons ('place cells') that selectively fire when an individual traverses a particular location within the environment [1,2]. One of the prominent properties of place cells is their tendency to 'remap' [3], i.e., form different representations for different environments by significantly changing their firing patterns. Remapping can be induced by changes in sensory input [3–8], motivational state [9–11], prior experience [12–14], or spatial trajectories [15–17].

While such alterations can induce remapping, it is well accepted that when an animal revisits an environment under the same experimental conditions the same spatial code is reused [18]. However, several studies have found instances in which place fields completely rearranged when the animal returned to the same environment after a short delay [19–23]. These observations raise the possibility that in certain cases either the

hippocampal spatial code is unstable or the stored memory is inaccessible, requiring the emergence of a new spatial representation [24]. An alternative hypothesis, however, is that multiple intact spatial representations of the same context can stably coexist in the hippocampus and be reliably retrieved from long-term memory. Differentiating between these alternative hypotheses requires comparing spatial representations across many consecutive visits to the same environment, and tracking such hippocampal representations over time-scales relevant for long-term memory.

Taking advantage of recent advancements in optical imaging [25] that enable tracking large populations of the same neurons over long time periods [26], we monitored hippocampal place cells of mice as they explored the same environment over many repeated visits, both within the same session, and across multiple sessions spanning weeks. Our results indicate that multiple representations of the same spatial context can coexist in the hippocampus and be stably retrieved over long time periods.

RESULTS

We performed large-scale Ca^{2+} imaging of pyramidal neurons in hippocampal area CA1 of freely behaving mice (n=27) during repeated explorations of the same familiar linear tracks (Figure 1A and B; see STAR methods) on 5-8 imaging sessions spanning 9-29 days (269-665 cells per session). Each imaging session consisted of five separate three-minute visits (trials) in the same environment, during which mice ran back and forth to collect water rewards at both ends of the track, with a three-minute inter-trial interval in which the mice were placed in an opaque bucket near the track (Figure 1C). Overall, for each mouse, we collected data from 25-40 trials within the same environment.

Distinct maps of the same familiar environment can coexist in hippocampal CA1

While in some mice place-cell specificity was maintained across all trials within a given session (Figure 1D left), most mice (19/27) exhibited a repertoire of distinct spatial representations (maps) across trials (Figure 1D right; compare trials 1, 2, and 4 with trials 3 and 5). The occurrence of multiple maps of the same environment was observed

in mice spanning different ages and genetic backgrounds (see STAR methods). Alternations between maps occurred coherently for the recorded population of place cells (Figure 1E and F, and Figure S1A). The overlap in the population of active cells across trials with different maps was lower than the overlap across trials with the same map (29.5±8.3% versus 49.3±7.5%, respectively), but higher than chance (Figure S1B and C). These results indicate that partially overlapping populations of neurons are active across different maps, consistent with previous studies that found shared ensemble activity patterns even across different spatial contexts [28,29,32]. Multiple maps for the same environment were observed using different cell-detection algorithms (Figure S1D-F), and were not explained by changes in the focal plane across trials (Figure S1G and H). The distribution of population vector (PV) correlations between trials was bimodal (Figure 1G), separating between pairs of trials with the same map and pairs with different maps. This bimodality was also observed at the level of specific spatial locations (Figure 1H) and individual cells (Figure 1I). We never found switching between maps within trials (0/874 three-minute trials), while across trials it occurred frequently (164/699 pairs of consecutive trials; Figure S2 and STAR methods), indicating that disconnection from the environment is required for the alternation to occur. Thus, the hippocampus can exhibit completely distinct spatial representations across separate visits to the same environment within a given imaging session.

Distinct maps are spatially informative and stable over weeks

Next, we used a probabilistic method for cell registration across sessions [26] to reliably track the same neurons over weeks (601-1,114 cells per mouse, estimated false-positive registration errors =5.6±2.8%, false-negative registration errors =4.1±1.8%; Figure S3). Similar to our observations within a given session, while some animals exhibited the same map across all trials from all days (Figure 2A left), the majority of mice (19/27) displayed multiple maps for the same environment, which recurred over weeks (Figure 2A right). We found no differences in cell registration accuracy between mice with one map and mice with multiple maps (Figure S3F). A pairwise similarity analysis across all trials from all sessions revealed a clear separation between the distinct maps (Figure 2B and C). We found no significant difference in mouse behavior

(fraction of stationary periods, number of track traversals, and running speed) between trials with different maps (Figure S4A-C), or between mice with one map and mice with multiple maps (Figure S5E-G). Thus, it does not seem that specific maps are used when mice are more motivated to run for reward, while others are used when mice are less motivated. There was also no difference in the effective number of maps for a given environment between mice that explored one environment and mice that explored two different environments (Figure S4D). Importantly, in all cases of multiple maps for the same environment, we found substantially high fractions of significantly tuned place cells and high levels of spatial information (Figure 2D-F), which were similar to those of mice that exhibited one map (Figure 2G and Figure S5A-D). Moreover, non-place cells displayed multiple activity patterns for the same environment, and switched between them in coherence with place cells (Figure 3), further reflecting the collective behavior of hippocampal neurons [27]. Taken together, these results indicate that multiple spatially informative representations of the same environment recur over long time periods.

We then studied the long-term stability of each map and its likelihood to be retrieved across days. First, we calculated the PV correlations between maps in different trials as a function of the elapsed time between them, and found that each map gradually changed over days (Figure 4A), consistent with previous studies [25,28]. Note, however, that even for elapsed times of two weeks, the PV correlations between trials in the same map (blue curve in Figure 4A) were still much higher than the PV correlations between trials in different maps, which consistently exhibited near-zero correlations (red curve in Figure 4A). By calculating the positional shifts across sessions, we found that even after ≥10 days, the majority of place cells that were active in a given map maintained their place fields when the same map recurred (Figure 4B; 60±7% versus 14±3% of cells exhibit ≤6 cm shift for the same map and across maps, respectively). While the effective number of maps did not change over the course of the experiment (Figure 4C, inset), the probability that the same map will recur gradually decreased as a function of the elapsed time (Figure 4C), indicating a drift in the prevalence of different maps over days. Importantly, the extent of the long-term changes in a given map was similar between mice with multiple maps and mice with one map (Figure 4D). These results

expand upon previous studies describing long-term changes in hippocampal ensemble activity [25,28,30,31], showing the contribution of two separate components to the observed dynamics: all-or-none switching between distinct maps with a gradually changing probability for a given map to recur, and gradual changes in the map itself.

Multiple maps are not mere rotations of the same spatial representation

We next sought to study in detail the relationships between the different maps of the same environment. To this end, we averaged the neuronal activity over all trials in each distinct map and calculated the PV correlations between all combinations of spatial locations, running directions, and maps (Figure 5A and B). While in some cases (15/40 map-pairs; see STAR methods) the maps were orthogonal (Figure 5A), in others they were not (Figure 5B). Specifically, in 24 out of 40 map-pairs, we found map segments that were similar under a flipped orientation of the environment (inverted order of the spatial locations and running directions in one of the maps; Figure 5B and C and Figure S6). For such flipped map-pairs, cells that encoded similar locations in one map tended to encode similar locations in the other map (Figure S6K), consistent with previous observations across different contexts [32]. While some mice exhibited higher PV correlations across different maps under a flipped orientation of the track (average PV correlation =0.07±0.11) compared to the same orientation (PV correlation =-0.01±0.06), these correlation values were still much lower than those observed across trials with the same map (PV correlation =0.58±0.14; Figure 5D).

Nonetheless, these analyses raise the possibility that in some of the mice the emergence of multiple maps for the same environment could be linked, at least partially, to 180-degree rotations in the representation of the linear track. Therefore, we tested whether multiple maps also occur in a setting that minimizes the possibility for disorientation between the two running directions. To this end, we made two modifications to our experimental setup (Figure 5E; see STAR methods): (1) During the inter-trial interval, we placed the mice in a transparent bucket instead of an opaque bucket, making the environment visible to the mouse throughout the session; (2) we used a set of proximal and distal visual cues designed to increase the polarization of the

environment (polarizing-cues experiment) instead of multiple cues scattered along the track and the surrounding enclosure (scattered-cues experiment shown in Figures 1-4 and Figure 5A-D). We imaged hippocampal activity of 10 additional mice that repeatedly explored the modified environments (8 sessions, 5 trials per session), and found that within a given session the same map occurred across all trials (Figure 5F and G). Specifically, the distribution of PV correlations across trials within the same session was unimodal (Figure 5G), unlike in the scattered-cues experiment (Figure 1G). These results further demonstrate that a spatial representation is persistent as long as the animal is not completely disconnected from the environment. However, comparing the maps across different sessions revealed that switching between maps still occurred in the majority of mice (6/10; Figure 5H and I). Importantly, unlike in the scattered-cues experiment, in the polarizing-cues experiment the maps were always orthogonal, not showing similarity under a flipped orientation of the track (Figure 5J and K). Overall, we found that multiple maps of the same environment can stably coexist even without signs of disorientation.

DISCUSSION

Here, we found that the population of hippocampal neurons can globally remap (i.e., change their firing rates and firing locations in an unstructured manner) within the same familiar environment without any apparent changes in sensory input or behavior. The fact that alternations between multiple representations of the same environment have been scarcely reported in the literature, but amply observed in this study, can be attributed to several differences between our experiments and previous studies. For example, while many studies in rodents used a single prominent visual orientation cue [3,4,6,7,33–35], we used multiple proximal and distal visual cues scattered along the track and the surrounding enclosure. The multiple cues may enable mice to utilize different cues as reference frames on different trials, leading to the formation of multiple maps. This interpretation is in line with the finding that hippocampal place cells can remap between different representations within the same environment when rodents rely on competing spatial reference frames [36–40]. Additionally, switching between maps required complete disconnection from the environment, and the probability of

using different maps increased with elapsed time between sessions. Thus, recording numerous separate visits to the same environment, combined with tracking the same neurons over weeks [25,26], increased the chances of observing multiple maps for the same environment, and allowed demonstrating their stable coexistence over long periods. Furthermore, avoiding pooling data from different trials enabled us to distinguish between cases in which place cells remap and cases wherein they maintain their place fields, attaining a more accurate description of place field stability and network dynamics (e.g., blue versus green curves in Figure 4A and B and Figure S1A-C). Overall, these results emphasize the importance of large-scale recordings and longitudinal analyses for interpreting neuronal activity.

What mechanism could support the coexistence of multiple maps for the same environment? The finding that remapping between representations is coherent across the recorded population (Figure 1E-I), involving both place and non-place cells (Figure 3), points to a network-level mechanism. Additionally, once a given map was retrieved, it persisted throughout the trial until the mouse was disconnected from the environment (Figure S2), allowing the network to reset and switch spatial representations between visits. This implies that the initial state of the network when entering the environment determines the retrieved map in a given visit. Finally, the convergence of the network to any map from a repertoire of spatial representations under the same experimental conditions is a hallmark of a non-linear dynamical system. Together, these results suggest the coexistence of different continuous attractors residing within the hippocampal network, consistent with the auto-associative properties attributed to the hippocampus [27,41–46].

In the framework of attractor neural networks, the structure of synaptic connectivity in the hippocampus defines an energy landscape [44,47], which determines the attractor states the network may converge to when encountering a novel environment (see illustration in Figure 6A and B). Low variability in the network's initial conditions or inputs between visits during familiarization could result in convergence to the same attractor on different instances, which might lead to the stabilization of a single attractor state,

corresponding to a single map (Figure 6A). Higher variability in initial conditions or inputs could result in convergence to different attractors on different visits, which may lead to the stabilization of separate attractor states, corresponding to multiple stable maps of the same spatial context (Figure 6B). Accordingly, manipulating the initial conditions of the network at the moment the mouse enters the environment may allow controlling the number or identity of the retrieved maps. Long-term changes in neuronal excitability or in synaptic connectivity [28,29,48] may alter the energy landscape of the network (Figure 6C). In line with our results, such changes should influence both the minimum points (the attractor states) and the separatrices (borders between basins of attraction) within the energy landscape, leading to gradual changes in the spatial representations (Figure 4A) [25,28] and in the probability of a given representation to recur (Figure 4C), respectively.

An example of initial conditions or inputs that could differ across separate encounters of an environment and influence its spatial representation is the orientation signal conveyed to the hippocampus through the encoding of head direction [49]. Indeed, we found evidence of a 180-degree rotation in the spatial representation of segments of the linear track in 60% of the map-pairs (Figure 5C), reflecting the two different possibilities for map realignment according to the geometry of the environment [50–54]. However, while in previous studies of map realignment, the majority of cells rotated coherently [50,52,54], in our experiment a relatively small portion of cells rotated across maps (14%; Figure S6L), with the rest exhibiting seemingly unstructured changes in their place fields. The hippocampal integration of a rotated head-direction signal with conflicting visual cues and other sensory inputs may have led to the remapping we observed. Consistently, inducing a mismatch between different hippocampal inputs can trigger global changes in hippocampal firing fields [20,35–37,55,56]. While realignment may explain the existence of two rotated maps for the same linear track, we often observed orthogonal maps (Figure 5A and C, Figure 5J and K, and Figure S6D-F) or more than two maps of the same linear track (Figure 2A right, Figure 2B and C top, Figure 2G, and Figure 5J). Furthermore, in the polarizing-cues experiment, mice had multiple maps for the same environment, but without any signs of 180-degree rotations

(Figure 5J and K). While we cannot rule out the possibility of other types of rotations, this is unlikely since for any coherent rotation, the activity would have been somewhat correlated between certain spatial locations across pairs of maps [54], and we did not find such correlations in our data. Thus, the initial state of the hippocampus or its inputs from other brain circuits, such as the entorhinal cortex [57,58], may have changed across trials, leading to the observed remapping.

While we have found multiple maps for the same environment in 60-70% of the mice, in the others, only one map was observed. One possibility is that there is a categorical difference between mice with one map and mice with multiple maps, which may reflect different strategies the mice employ to perform the task. However, we did not find differences in behavior between mice with one map and mice with multiple maps to support this notion. Alternatively, different mice may represent samples from a continuous distribution of remapping probability, with one-map mice at the lower end of the distribution. If spatial representations become more stable with familiarization, it could be that for mice with an initially low remapping probability, only one attractor is stabilized in the first few visits to the environment, further decreasing the probability of observing a different map to the point it is practically zero. Thus, a continuous distribution of remapping probability in a novel environment could be transformed into a categorical difference in a familiar environment.

Previous work has demonstrated that in rats with age-related memory impairment [22,23], blocked NMDA receptors [59], or blocked protein synthesis [60], place fields could be rearranged when the animals returned to the same environment after a short delay. These observations were attributed to cases where memory is impaired and the contextual information is inaccessible. We have found, however, that distinct maps stably recur over multiple weeks and are spatially informative, demonstrating that they are accessible and can be reliably stored and retrieved from long-term memory. Since redundancy may increase the robustness of memory retrieval [63], storing multiple representations of the same environment may even confer a functional advantage. Furthermore, recent work in artificial intelligence showed that storing multiple maps of

the same place helps flexible navigation in a non-stationary environment [64,65]. Considered alongside our results, these findings suggest that having multiple maps of the same environment could be advantageous under certain real-life situations, such as navigating in a dynamic environment.

A related open question is whether distinct maps of the same environment are cognitively interlinked or, alternatively, represent separate perceptions or memory traces. Previous studies have shown that performance in a hippocampal-dependent task can be preserved despite global remapping [61,62]. Consistently, we did not find behavioral differences across trials in different maps (Figure S4A-C), or between mice with one map and mice with multiple maps for the same environment (Figure S5E-G). A possible explanation for this result is the existence of a “universal map”, linking between different representations, outside the hippocampus [5,20]. Contrary to this view, silencing place cells that are active in a familiar environment can immediately trigger the emergence of an alternative spatial representation and abolishment of the behavior previously associated with that environment [24]. Finding multiple maps for the same environment while animals perform a hippocampal-dependent memory task could help determine whether the same environment can be perceived as multiple unrelated contexts, or whether the memory associated with a given environment can be supported by multiple neuronal representations.

AKNOWLEGMENTS

Y.Z. is a CIFAR-Azrieli Global Scholar in the Brain, Mind & Consciousness program, and the incumbent of the Daniel E. Koshland Sr. Career Development Chair. Y.Z. is supported by grants from the Abraham and Sonia Rochlin Foundation, the Hymen T. Milgrom Trust, the Israel Science Foundation (grant 2113/19), the Human Frontier Science Program, and the European Research Council (ERC-StG 638644). We thank Timothy O'Leary, Christoph Schmidt-Hieber, Ofer Yizhar, Nachum Ulanovsky, Michal Rivlin, Tamir Eliav, Arseny Finkelstein, Yitzhak Norman, Dori Derdikman, Noa Sadeh, Linor Balilti-Turgeman, Meytar Zemer, Eyal Bitton, Daniel Dietch, and Natalie Page for helpful advice and comments on the manuscript.

AUTHOR CONTRIBUTIONS

Conceptualization, A.R; Methodology, L.S., N.G., A.R., and Y.Z.; Software & Formal Analysis, L.S.; Investigation, N.G., Y.R., and H.B.; Writing, L.S., N.G., A.R., and Y.Z.; – Review & Editing, L.S., N.G., H.B., Y.R., A.R., and Y.Z.; Supervision, A.R. and Y.Z.

DECLARATION OF INTERESTS

The authors declare no competing interests

Figure 1. Distinct spatial representations (maps) of the same familiar environment can coexist in hippocampal CA1

(A-C) Experimental setup. (A) Ca^{2+} imaging using a miniaturized head-mounted microscope in hippocampal CA1 of freely behaving mice. (B) Mice repeatedly explored familiar environments to collect water rewards at the two ends of a linear track. (C) Experiment timeline: mice explored the same environments in 5-8 imaging sessions spanning weeks, each comprising five three-minute trials. During the three-minute inter-trial interval (ITI), mice were placed in an opaque bucket near the track. (D) Five example place cells in a mouse with one map (left) and a mouse that switches between different maps across trials (right; compare trials 1, 2, and 4 with trials 3 and 5). Top: Mouse position during neuronal activity events (red and green dots for rightward and leftward running directions, respectively) overlaid on the mouse trajectory (blue curve). Black horizontal lines show separation between trials. Bottom: Corresponding spatial activity shown separately for each map (blue and magenta dashed curves) and combined (solid black curve). (E) Rate maps for all significant spatially modulated cells for the same trials shown in D. Note that the population of place cells coherently switches between distinct maps across five consecutive trials (right). The cells are sorted according to their preferred position along the track in the trial marked by a green frame, i.e., the first and fifth trials in the top and bottom rows, respectively. (F) PV correlations between the representations of different spatial locations across trials reveal two distinct maps. In E and F, data were pooled from both running directions. (G) The distribution of average PV correlations across trials (gray bars) is bimodal, unlike the shuffled data (black curve). (H and I) Remapping is coherent across all positions and cells. Distribution of PV correlations between the neuronal representation of the same locations (H), and rate map correlations between the same cells (I) across trials are higher for the same map (blue) compared to different maps (red) or shuffled data (black). Shuffled data was obtained by shuffling the identity of the cells (G and H) or spatial locations (I) across trials. See STAR methods for the procedure used for the identification of the different maps. Data in G-I was pooled from 19 mice with multiple maps for the same environment. For mice that explored two environments (7/19), data from each environment was analyzed separately. See also Figures S1 and S2.

Figure 2. Distinct maps of the same environment are informative about space
(A) PV correlations between the representations of different locations (pooled from both running directions) across trials from three consecutive imaging sessions for a mouse with one map (left) and a mouse with multiple maps (right). (B-D) One example mouse exhibiting three maps (top) and another mouse exhibiting two maps (bottom) for the same environment . (B) Average PV correlations between all trials from all sessions in a given mouse show that different maps recur over weeks. (C) Projection of the correlation matrix in B on its two leading eigenvectors reveals separate clusters, each corresponding to a distinct map. This clustering was used to define the different maps and identify the trials in which each of them was used. Inset: correlation matrix in B, ordered according to the maps' clustering. (D) Average spatial information per neuron and percentage of place cells in each trial are similar across distinct maps (colored dots), and are significantly higher than chance (black dots). (E, F) Similar percentage of place cells (E) and average spatial information (F) are observed across the different maps (blue) in 25/26 and 24/26 cases, respectively ($p>0.05$, one-way two-sided ANOVA). For mice that explored two environments (9/27), each environment was analyzed separately. Each data point is a comparison between all the trials from two distinct maps of the same environment in the same mouse (mean±SD). For mice with more than 2 maps, all map-pairs are shown. The percentage of place cells and average spatial information were higher than chance (black) for all maps. Data in E and F was pooled from 19 mice with multiple maps for the same environment. (G) The distribution of the effective number of maps for one-map (cyan) and multiple-map mice (magenta). The effective number of maps is calculated based on the occurrence probability of the different maps (see STAR methods). For mice that explored two environments, the effective number of maps was averaged across environments. See also Figures S3-S5.

Figure 3. Place and non-place cells coherently switch between multiple representations of the same environment

(A and B) The average rate correlations between all trials from seven recording sessions for place cells (A) and non-place cells (B) show coherent alternations between distinct representations over multiple days. The average rate is based on the overall activity of each cell in a given trial, irrespective of mouse position. Place cells are cells that were significantly tuned to position in at least one of the two trials for which the rate correlation was computed. (C) Rate correlations for non-place cells across pairs of trials within the same session increase with the rate correlations for place cells across the same pairs of trials. Inset: same analysis for pairs of trials across different sessions. Data in C was pooled from 19 mice with multiple maps for the same environment.

Figure 4. Distinct maps of the same environment are stable over weeks
(A) Average PV correlation between trials as a function of elapsed time between them (mean±SD) are higher for the same map (blue) than across different maps (red; two-sample one-sided t-test, $p<0.005$ for all elapsed times). Average across all trials is shown in green. For mice that explored two environments, the PV correlation was averaged across environments. (B) Distribution (mean±SEM) of place-field positional shifts for pairs of trials ≥10 days apart in the same map (blue), different maps (red), and overall (green), showing place-field stability over long time periods within the same map. Inset: cell fractions with ≤6 cm positional shifts across trials ≥10 days apart. (C) The probability of the same map to occur on two given trials as a function of elapsed time between them (mean±SEM) shows a gradual change in the recurrence of maps over days. Inset: effective number of maps does not differ between the first and second half of the experiment (1.69±0.72 versus 1.84±0.71, mean±SD for sessions 1-4 and sessions 5-8, respectively; matched-pairs two-sided $t\text{-test}_{(18)}$, $t=-1.43$, $p=0.17$). Data in A-C was pooled from 19 mice with multiple maps for the same environment. (D) Similar PV correlations (mean±SD) between trials within the same map are observed for one-map (cyan) and multiple-map (magenta; two-sided t-test, $p>0.05$ for all elapsed times) mice. For mice that explored two environments, the PV correlation was averaged across environments. Data pooled from 19 mice with multiple maps and 8 mice with one map for the same environment.

Figure 5. The occurrence of multiple maps for the same environment is not explained by mere rotations of the spatial representation

(A and B) PV correlations between all combinations of spatial locations, running directions, and maps reveal in some cases a similarity across representations of different locations in the environment. PVs were obtained by averaging over all trials in a given map. (A) Example with no relationship observed between the different maps of the same track. (B) Example of a similar representation in one segment of the environment under a flipped track orientation, i.e., the representation of the beginning of rightward running in Map 1 resembles that of the beginning of leftward running in Map 2. The only case in which we found a coherent rotation of the entire representation across two maps is shown in Figure S6C. (C) Average PV correlation between map pairs of the same environment under a flipped track orientation versus same orientation. The examples shown in A and B (mice 6 and 2) are indicated by colored markers. For mice with more than two maps, all map-pairs are shown. (D) Higher PV correlations are observed across trials from different maps under a flipped environment orientation (middle) than under the same orientation (top), but are far lower than those observed across trials with the same map (bottom). Only trial-pairs from the same day are presented. Data in C and D were pooled from 19 mice with multiple maps for the same environment. (E-K) Polarizing-cues experiment. (E) Mice repeatedly explored modified linear environments with polarizing distal and proximal visual cues. A transparent bucket was used for the ITI. (F) PV correlations between the representations of different locations (pooled from both running directions) across trials from three consecutive imaging sessions for a mouse that explored the modified environment. (G) Distribution of average PV correlations across trial pairs within the same session (gray bars) is unimodal and markedly higher than shuffled data (black curve; obtained by shuffling the identity of the cells across trials). Data was pooled from 10 mice, each exploring two environments on eight imaging sessions. (H) The distribution of average PV correlations across session pairs (gray bars) is bimodal, unlike that of shuffled data (black curve; obtained by shuffling the identity of the cells across sessions). For each session, the PVs were averaged across all five trials. Data was pooled from six mice that had multiple maps for the same environment. (I) Distribution of the effective number of maps, averaged across environments, for one-map (cyan) and multiple-map (magenta) mice. (J) PV correlations between all combinations of spatial locations, running directions, and maps show no relations across the different representations of the environment. (K) Average PV correlation between map pairs of the same environment under a flipped track orientation versus same orientation shows no relations across maps. For mice with more than two maps, all map pairs are shown. See also Figure S6.

Figure 6. The auto-associative properties of the hippocampus provide a possible mechanism for the formation and evolution of multiple maps for the same environment
(A and B) Illustration of the formation of either one (A) or two (B) stable attractors in the energy landscape, corresponding to one or two stable maps of the same environment. While under low variability in initial conditions or inputs (dense arrows), the network might always converge to the same attractor (black asterisk), hence stabilizing only one map (A), under higher variability in initial conditions or inputs (sparse arrows), it can converge to either of two different attractors (two black asterisks), thus stabilizing two maps (B). (C) Illustration of the energy landscape supporting two stable maps of the same familiar environment as it gradually changes over long time scales (blue and red curves show the energy landscape of the same network on different days). The drift in the attractors and separatrix leads to changes in the maps and the tendency to converge to each map, respectively. Hence, small changes in the energy landscape can cause the network to converge to different attractors (blue and red asterisks) under the same initial conditions or inputs.

STAR * METHODS

LEAD CONTACT AND MATERIALS AVAILABILITY

This study did not generate new unique reagents. Further information and requests for resources and reagents should be directed to the Lead Contact, Yaniv Ziv (yaniv.ziv@weizmann.ac.il).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Animals

All procedures were approved by the Weizmann Institute IACUC. A total of 37 male mice, aged 3-10 months at the start of imaging, were used in this study, out of which, data from 5 mice was taken from a previous publication [28]. Out of these 37 mice, 27 performed the scattered-cues experiment (Figures 1-4 and Figure 5A-D), and 10 performed the polarizing-cues experiment (Figure 5E-K). Twenty six mice were virus injected C57BL/6 (3 injected with AAV2/5-CaMKIIa-GCaMP6s and 23 injected with AAV2/5-CaMKIIa-GCaMP6f), and 11 were Thy1-GCaMP6f transgenic mice (Jackson stock number 025393; 6 with a C57BL/6 background and 5 with an ICR(CD-1) background) [66] (see table below). Mice were housed with 1-4 cage-mates in cages with running wheels in a 12:12 hour reverse light cycle (lights off 10am-10pm).

	AAV2/5-CaMKIIa-GCaMP6f	AAV2/5-CaMKIIa-GCaMP6s	Thy1-GCaMP6f
Scattered cues	10/13	3/3	6/11
Polarizing cues	6/10	-	-

Number of mice with multiple maps for the same environment out of the total number of mice across the different genetic backgrounds and experiments

METHOD DETAILS

Surgical procedures

For the virus-injected C57BL/6 mice, we first injected into the CA1, under isoflurane anesthesia (1.5-2% volume), 450 nL of viral vector AAV2/5-CaMKIIa-GCaMP6s or AAV2/5-CaMKIIa-GCaMP6f (~$2X10^{13}$ particles per ml, packed by University of North Carolina Vector Core)[67] at stereotatic coordinates: -1.9 mm anterio-posterior, -1.4 mm mediolateral, -1.6 mm dorsoventral from Bregma. Mice were then allowed to recover for at least two weeks. We then implanted all mice with a glass guide tube directly above the

CA1. We used a trephine drill to remove a circular part of the skull centered posteriolateral to the viral injection site. We removed the dura and cortex above the CA1 by suction with a 29 gauge blunt needle while constantly washing the exposed tissue with sterile PBS. We then implanted an optical guide tube with its window just dorsal to, but not within, area CA1 and sealed the space between the skull and the guide tube using 1.5% agarose in PBS. The exposed skull areas were then sealed with Metabond (Parkell, Edgewood, NY) and dental acrylic.

Ca^{2+} imaging and behavioral setup

To perform time-lapse imaging in freely behaving mice using an integrated miniature fluorescence microscope (nVistaHD, Inscopix), we followed a previously established protocol [25]. Mice were water-restricted and maintained above 80% of their initial weight. We habituated the mice to human handling by allowing them to walk on the experimenters' hands. At least three weeks after guide-tube implantation, mice were imaged under isoflurane anesthesia using a two-photon microscope (Ultima IV, Bruker, Germany). We inserted a microendoscope consisting of a single gradient refractive index lens (1 mm diameter, Inscopix) into the guide tube and examined Ca^{2+} indicator expression and tissue health. We selected for further imaging only those mice that exhibited homogenous GCaMP6 expression and appeared to have healthy tissue. For the selected mice, we then affixed the microendoscope within the guide tube using an ultraviolet-curing adhesive (Norland, NOA81, Edmund Optics, Barrington, NJ). Next, we attached the microscope's base plate to the dental acrylic cap using a light-cured adhesive (Flow-It ALC, Pentron, Orange, CA). To record mouse behavior, we used an overhead camera (DFK 33G445, The Imaging Source, Germany), which we synchronized with the integrated microscope. Ca^{2+} imaging was performed at 20 Hz.

Scattered-cues experiment

The data set from this experiment is presented in Figures 1-4 and Figure 5A-D. Before starting the experiment, we pre-trained the mice to run back and forth to collect water rewards on one or two linear tracks (Environments A and B). In each environment, we placed the track in a square enclosure within the recording room. The environments

differed in their geometry and had distinct sets of visual and tactile cues, overhead lights, flavored liquid rewards, and odor cues. Environment A was a straight 96-cm-long track, while Environment B was an L-shaped track consisting of two 48-cm-long arms. Each environment had a separate set of curtains with different colors and patterns, which made up the walls of the recording enclosure (distal cues), and also had distinct visual cues scattered along the walls of each track (proximal cues). Before the beginning of each session, we wiped the tracks with differently scented paper towels (0.5% acetic acid for environment A and 10% ethanol for environment B). The water reward was dispensed at both ends of the tracks using a custom-made computer-controlled device. Mice were pre-trained over 6–11 days while carrying the head-mounted microscope, until they ran at least 60 times the entire length of each track in two consecutive days. Pre-training and imaging sessions consisted of five 3-minute-long trials, with an inter-trial interval (ITI) of 3 minutes, during which the mice were put in an opaque circular plastic bucket, which was placed at a specific fixed position on top of the track. When necessary, the bucket was rotated during the ITI to untangle the coiling of the microscope data wire. Special care was taken to ensure each mouse was placed on the track in the same way across all trials. We imaged the mice every 2 or 4 days for a total of 5-8 recording days. Nine mice explored both environments on each day in two sessions that were separated by 4–5 hours, and 18 mice explored only environment A. On each recording day, we imaged 1-7 mice.

Polarizing-cues experiment

The data set from this experiment is presented in Figure 5E-K. We designed this experiment to test whether multiple maps also occur in a setting that minimizes the possibility for disorientation between the two running directions of the linear tracks. To this end, we created environments with highly polarizing orientation cues. In polarized environment A (straight track), three of the enclosure's walls were black and one white. In polarized environment B (L-shaped track), two of the enclosure's outer walls were black, one was white, and one featured vertical black and white stripes. In both environments, the proximal visual cues covered only half the length of the track (using a unique cue-set for each environment). In polarized environment A, we created grooves

in the track's floor (tactile cues) that were parallel to the running direction of the mice on one side of the track length and perpendicular to the running direction on the other side. Before beginning Ca^{2+} imaging, we pre-trained the mice for four days in a 56-cm-long linear track with a different set of visual and olfactory cues. Imaging sessions had the same structure as described above for the scatted cues experiment. During the 3-minute ITI, the mice were placed in a transparent circular plastic bucket instead of an opaque one, so that the environment remained visible throughout the ITI. In this experiment, the bucket was never rotated during the ITI, to further avoid disorienting the mice. We imaged the mice every three days for a total of eight recording days. All 10 mice explored both environments in each day of the experiment in two sessions that were separated by 4–5 hours. On each recording day, we imaged 5 mice.

Mouse position tracking

We analyzed mouse behavior videos using a custom MATLAB (Mathworks) routine that detects the mouse's center of mass in each frame. We then used the estimated position to calculate its velocity, and applied a smoothing filter (rectangular window of 250 msec) to the calculated velocity.

Processing Ca^{2+} imaging data

We processed the Ca^{2+} imaging data using commercial software (Mosaic, version 1.1.1b, Inscopix) and custom MATLAB routines as previously described [25,28]. To increase the computation speed, we spatially down-sampled the data by a factor of four in each dimension (final pixel size of 2.3 X 2.3 μm). To correct for non-uniform illumination, both in space and in time, we normalized the images by dividing each pixel by the corresponding value from a smoothed image. The smoothed image was obtained by applying a Gaussian filter with a radius of 100 μm to the movies. Normalization also enhanced the appearance of the blood vessels, which were later used as stationary fiducial markers for image registration. We used a rigid-body registration to correct for lateral displacements of the brain. This procedure was performed on a high-contrast sub-region of the normalized movies in which the blood vessels were most prominent. The movies were transformed into relative fluorescence changes, $\frac{\Delta F(t)}{F_0} = (F(t) - F_0)/F_0$,

where F_0 is the value for each pixel averaged over the entire recording duration. For cell detection, the movies were down-sampled in time by a factor of five. We detected spatial footprints corresponding to individual cells using an established cell-detection algorithm that applies principal and independent component analyses (PCA and ICA) [68]. For each spatial footprint, we used a threshold of 50% of the footprint's maximum intensity, and each pixel that did not cross this threshold was set to zero. After the cells were detected, further cell sorting was performed to identify the spatial footprints that follow a typical cellular structure. This was done by measuring the footprint area and circularity, discarding those whose radius was smaller than 5 μm or larger than 15 μm or which had a circularity smaller than 0.8. In some cases, the output of the PCA/ICA algorithm included more than one component that corresponded to a single cell. To eliminate such occurrences, we examined all pairs of cells with centroid distances <18 μm; whenever their traces had a correlation >0.9, the cell with the lower average event peak amplitude was discarded. We detected 269-665 per imaging session.

In the analysis presented in Figure S1D-F, calcium dynamics were also extracted using CNMF-e [69], an extension of the constrained non-negative matrix factorization method [70], for one-photon microendoscopic data.

Detection of Ca^{2+} events

Ca^{2+} activity was extracted by applying the thresholded spatial footprints to the CA^{2+} movies at full temporal resolution (20 Hz) $\Delta F(t)/F_0$. Baseline fluctuations were removed by subtracting the median trace (20 sec sliding window). The Ca^{2+} traces were smoothed with a low-pass filter with a cutoff frequency of 2 Hz. Ca^{2+} candidate events were detected whenever the amplitude crossed a threshold of 4 or 5 median absolute deviations (MAD) for GCaMP6s or GCaMP6f, respectively. We considered for further analysis only candidate Ca^{2+} events with an indicator decay time for GCaMP6s or GCaMP6f equal to or longer than 600 msec or 200 msec, respectively. In order to avoid the detection of several peaks corresponding to a single Ca^{2+} event, only peaks 4 or 5 MAD higher than the previous peak (within the same candidate event) and 2 or 2.5 MAD higher than the next peak for GCaMP6s or GCaMP6f, respectively, were regarded as true events. We

set the Ca^{2+} event occurrence to the time of the peak fluorescence. To mitigate the effects of crosstalk (i.e., spillover of Ca^{2+} fluorescence from neighboring cells), we adopted a conservative approach, allowing only one cell from a group of neighbors (pairs of cells with centroid distances <18 μm) to register a Ca^{2+} event in any 200 msec time window. If multiple Ca^{2+} events occurred within ~200 msec in neighboring cells, we retained only the event with the highest peak $\Delta F(t)/F_0$ value. If two neighboring cells had a correlation >0.9 in their events, the cell with the lower average peak amplitude was discarded. After the events were identified, further event sorting was performed to find the cells with sufficient signal-to-noise ratios. This was done by measuring the event rate and the average event peak amplitude for each cell and discarding those whose event rate was smaller than 0.01 Hz or which had an average event amplitude smaller than 1% ($\Delta F(t)/F_0$).

Registration of cells across sessions

To identify the same neurons across multiple imaging sessions, we used a probabilistic method for cell registration [26]. This method models the distribution of spatial correlations and centroid distances for neighboring cells from different recording sessions (candidates for being the same cell) as a weighted sum of the distributions of two subpopulations: same cells and different cells (Figure S3D). Then, based on the model that best fits the data, the method estimates the probability of each candidate in the data set to be the same cell (P_{same}). This allows estimating the overall rates of false-positive errors (different cells falsely registered as the same cell) and false-negative errors (the same cell falsely registered as different cells), providing a P_{same} registration threshold that is optimized to the data set of each mouse. The threshold used for registration controls the tradeoff between false-positive and false-negative errors. Therefore, we chose a registration threshold of P_{same}=0.5, which constitutes an appropriate balance between false-positive and false-negative registration errors [26]. Note that while some of the same-cell candidates had intermediate spatial correlation values, the vast majority of such cell pairs exhibited either very high spatial correlations (suggesting they are the same cell) or very low spatial correlations (suggesting they are different cells; Figure S3D). Overall, we could track the same neurons across multiple imaging days with estimated false-positive registration errors =5.6±2.8% and false-negative registration errors =4.1±1.8% (Figure

S3E). Furthermore, in cases with multiple candidates that cross the registration threshold, only the pair with the highest P_{same} was registered as the same cell, thus avoiding some of the false positive errors, the result of which was lower percentages of false-positive errors than those estimated by the probabilistic model. Importantly, since the microscope remained head-mounted throughout each session (i.e., maintaining exactly the same focal plane and field of view), the processing of the data and the detection of cells was done together for all trials within each session. Therefore, no cell registration was required across trials from the same session.

Rate maps

For place field analysis, we considered periods wherein the mouse ran >1 cm/sec. We divided each track into 24 bins and computed the time spent in each bin and the number of Ca^{2+} events per bin. We smoothed these two maps ('occupancy' and 'Ca^{2+} event number') using a truncated Gaussian kernel (σ=1.5 bins, size=5 bins) [25,28]. We then computed the activity map (event rate per bin, excluding the last 2 bins at each end of the tracks) for each neuron by dividing the smoothed map of Ca^{2+} event numbers by the smoothed map of occupancy. We separately considered place fields for the two running directions on the linear track. For each running direction, we defined the preferred position to be the position with a peak value within the activity map.

Place tuning analysis

For this analysis, we focused only on active cells (≥3 detected events in a given trial). We computed the spatial information (in bits per event) of these cells using the unsmoothed-event-rate map of each cell, as previously described [71]:

$$Spatial\ information = \sum_i p_i \left(r_i / \bar{r}\right) log_2 \left(r_i / \bar{r}\right),$$

where p_i is the probability of the mouse to be in the i^{th} bin (time spent in i^{th} bin/total running time); r_i is the Ca^{2+} event rate in the i^{th} bin; and $\bar{r}$ is the overall mean Ca^{2+} event rate. We then performed 1,000 distinct shuffles of animal locations during Ca^{2+} events, accounting for the spatial coverage statistics for the relevant trial and running direction, and calculated the spatial information for each shuffle. This yielded the p-value of the

calculated spatial information relative to the shuffles. Cells with spatial information higher than that of 95% of their shuffles were considered significant place cells. The fraction of place cells in each trial was defined as the number of significant cells out of the number of active cells in that trial.

Population vector correlation

To determine the level of similarity between representations of the environment on different trials, we calculated the mean population vector correlation between them [7]. For each spatial bin (excluding the last 2 bins at each end of the tracks, and separating between right and left running directions), we defined the population vector as the mean event rate for each place cell (cells that were significantly spatially tuned in at least one of the two trails) given that bin's occupancy. We computed the correlation between the population vector (PV correlation) in one trial with that of the matching location in another trial, and averaged the correlations over all positions. In addition, to study the relationships between representations, we averaged the rate maps over all trials in a given map. We then calculated the PV correlations between all the combinations of spatial locations, running directions, and distinct maps. To measure the similarity between pairs of maps under a flipped orientation of the environment, we inverted in one of the maps the order of the spatial bins and switched between left and right running directions. To determine whether distinct maps of the same environment are orthogonal or show a relationship, we calculated the distribution of PV correlations for all spatial locations using a shuffled version of the data. We then counted for each pair of maps of the same environment the fraction of spatial locations that were significantly correlated (higher than 95% of all shuffles). Finally, we asked whether the number of spatial locations that were significantly correlated across maps have <0.05 probability to be obtained by chance. In practice, 6 correlated spatial locations out of a total of 40 spatial locations (across the two running directions) were enough for a pair of maps to be considered not orthogonal (using Bonferroni correction for the two possible orientations of the track). For a pair of maps to be considered related specifically under a flipped orientation of the track, 5 correlated spatial locations out of a total of 40 spatial locations were enough (no need in this case for the Bonferroni correction for the two possible orientations).

Rate map correlation

For each place cell, we determined the level of similarity between the representations of the environment on different trials by calculating the rate map correlation (Pearson correlation) between the rate maps in two trials (done for each running direction separately). To measure the similarity between spatial tuning under a flipped orientation of the environment, we inverted in one of the maps the order of the spatial bins and switched between left and right running directions. Cells that were significantly spatially tuned in at least one of the two trails were considered to be place cells.

Distinction between different spatial representations (maps)

To distinguish between the different maps, we calculated the PV correlations across all trials recorded from the same mouse. To cluster the trials into the different maps, we computed the eigenvectors of the PV correlation matrix. Then, we projected the columns of the PV correlation matrix (each corresponding to a different trial) onto its two leading eigenvectors and clustered the trials in the dimensionality reduced PV correlations space. Since there were rarely cases of ambiguity in the clustering of different trials, determining the number of clusters and the assignment of each trial to a specific cluster was performed manually. We validated the manual assignment by applying two different unsupervised methods to the leading five dimensions of the reduced PV correlations space: K-means and topological clustering [72]. These validations yielded 96.2% and 98.5% agreement with the manual assignment for K-means and topological clustering, respectively. Only maps that occurred more than once were considered as one of the alternative representations of a given environment, and those that occurred only once were discarded from the analysis.

Effective number of maps

To quantify the tendency of a mouse to represent a given environment with multiple maps, we measured the number of trials in which each map occurred. We then calculated the effective number of maps based on the maps' entropy (H_{maps}), accounting for the probability of each map to occur in any given trial:

$$N_{effective} = 2^{H_{maps}} = 2^{-\sum_i P_i \log_2 P_i} ,$$

where P_i is the fraction of trials in which the i^{th} map occurred. This definition aligns with the naïve definition of the number of maps in the case of a uniform distribution (e.g., for 50% of the trials in each of two maps, $N_{effective}$=2; for 33% of the trials in each of three maps, $N_{effective}$=3), and may result in a non-integer in the case of a non-uniform distribution (e.g., for 25% of the trials in map 1 and 75% in map 2, $N_{effective}$=1.48).

Map stability and recurrence probability

To quantify the long-term dynamics of a given map, we calculated the average PV correlation between pairs of trials as a function of the elapsed time between them.

To quantify the long-term dynamics of map recurrence probability, we calculated the probability that the same map recur on two different trials as a function of elapsed time.

Average rate correlation

To quantify the similarities in activity patterns between trials, without accounting for place tuning, we measured the overall event rate for each neuron in each trial (irrespective of mouse position). We then calculated the rate correlation as the Pearson correlation between the vectors (for all neurons) of event rates. This calculation was also performed separately for place and non-place cells. Cells that were significantly spatially tuned in at least one of the two trails were considered to be place cells.

Overlap

To quantify the overlap between two trials, we counted the number of cells with ≥5 events in both trials and then divided the resultant figure by the total number of cells with ≥5 events in at least one of the trials.

Within-trial stability

To quantify the within-trial stability of a given map, we measured map similarity at a higher temporal resolution than the full trial length. First, we calculated the average PV correlations between two contiguous (non-overlapping) running windows of 60 sec within each imaging session (consisting of five trials). Both time windows were shifted

together throughout the entire session. Time points for which the PV correlation significantly decreased indicated a map switch (Figure S2A). We also calculated the average correlations between the PVs of a 60-second running window and a stationary reference window representing the average PVs from all trials in a given map, and repeated this calculation for all distinct maps. This analysis showed that at time points for which the PV correlations with one map significantly decreased, the PV correlations with another map increased (Figure S2B). Based on this analysis, we devised a map switch detector. Whenever there was a switch in the identity of the reference map with which the running window had the highest PV correlation (provided that the PV correlation was >0.2), a map switch was assigned to that time point. Then, we counted the number of map switches that occurred across trials (during the 30 seconds after the beginning or before the end of a trial), and the number of switches that occurred within a trial (during the middle two minutes of a trial). Although the defined within-trial time interval was two times larger than the across-trials time interval, we never detected switching between maps within trials (0 out of a total 874 three-minute trials), while across trials it was detected frequently (164 out of a total of 699 pairs of consecutive trials; Figure S2). Additionally, we calculated the average PV correlations between the first half and the second half of the same trials. To measure within-trial stability at an even finer temporal resolution, we also calculated the correlations between the PV of a single frame (time bin of 50 msec) and a stationary reference window representing the average PVs from all trials in the same map or in another map. We then compared between the distributions of single frame correlations to the same map reference and to other maps references.

Behavioral analysis

To compare mice behavior across the different maps, we measured the percentage of stationary periods, track traversal, and the average speed of the mouse. The percentage of stationary periods was defined as the amount of time the mice ran at <1 cm/sec out of the total time of each trial. The number of track traversals was measured based on the number of times the mouse reached either of the two track ends. The average speed was computed based on the periods for which the mouse was not stationary (≥1cm/sec). We

then divided the trials according to the different maps used in each trial and compared between the maps. We also compared the behavior between mice with one map and mice with multiple maps for the same environment.

QUANTIFICATION AND STATISTICAL ANALYSIS

We analyzed unrelated samples by way of a two-sided t-test or one-way two-sided analysis of variance (ANOVA). Bonferroni correction was performed when conducting multiple comparisons. For mice that explored two different environments (9/27 in the scattered-cues experiment and 10/10 in the polarizing-cues experiment), we averaged the data for each mouse over the two environments. Wilcoxon signed-rank test was used for comparing the effective number of maps between mice that explored one environment and mice that explored two environments. All statistical details, including the specific statistical tests, the values of n, and what it represents are specified in the corresponding figure legends. All statistical analyses were conducted using MATLAB (Mathworks).

DATA AND CODE AVAILIABILITY

The data set and code supporting the current study have not been deposited in a public repository because it is analyzed for an additional work that is still in progress, but are available from the corresponding author upon reasonable request.

REFERENCES

1. O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. Brain Res. *34*, 171–175.
2. O'Keefe, J., and Nadel, L. (1978). The hippocampus as a cognitive map. (Oxford: Clarendon Press.).
3. Muller, R.U., and Kubie, J.L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J. Neurosci. *7*, 1951–1968.
4. Fyhn, M., Hafting, T., Treves, A., Moser, M.B., and Moser, E.I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. Nature *446*, 190–194.
5. Geva-Sagiv, M., Romani, S., Las, L., and Ulanovsky, N. (2016). Hippocampal global remapping for different sensory modalities in flying bats. Nat. Neurosci. *19*, 952–958.
6. Lever, C., Wills, T., Cacucci, F., Burgess, N., and O'Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. Nature *416*, 90–94.
7. Leutgeb, S., Leutgeb, J.K., Barnes, C.A., Moser, E.I., McNaughton, B.L., and Moser, M.B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. Science *309*, 619–623.
8. Terrazas, A., Krause, M., Lipa, P., Gothard, K. M., Barnes, C. A., & McNaughton, B.L. (2005). Self-motion and the hippocampal spatial metric. J. Neurosci. *25*, 8085–8096.
9. Markus, E.J., Qin, Y., Leonard, B., Skaggs, W., McNaughton, B., and Barnes, C. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. J. Neurosci. *15*, 7079–7094.
10. Kennedy, P.J., and Shapiro, M.L. (2009). Motivational states activate distinct hippocampal representations to guide goal-directed behaviors. Proc. Natl. Acad. Sci. *106*, 10805–10810.
11. Moita, M. A., Rosis, S., Zhou, Y., LeDoux, J. E., & Blair, H.T. (2004). Putting fear in its place: remapping of hippocampal place cells during fear conditioning. J.

Neurosci. *24*, 7015–7023.

12. Quirk, G.J., Muller, R.U., and Kubie, J.L. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. J. Neurosci. *10*, 2008–2017.
13. Tanila, H. (1999). Hippocampal place cells can develop distinct representations of two visually identical environments. Hippocampus *9*, 235–246.
14. Bostock, E., Muller, R.U., and Kubie, J.L. (1991). Experience-dependent modifications of hippocampal place cell firing. Hippocampus *1*, 193–205.
15. Wood, E.R., Dudchenko, P.A., Robitsek, R.J., and Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron *27*, 623–633.
16. Ferbinteanu, J., and Shapiro, M.L. (2003). Prospective and retrospective memory coding in the hippocampus. Neuron *40*, 1227–1239.
17. Frank, L.M., Brown, E.N., and Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. Neuron *27*, 169–178.
18. Colgin, L.L., Moser, E.I., and Moser, M.B. (2008). Understanding memory through hippocampal remapping. Trends Neurosci. *31*, 469–477.
19. Kentros, C.G., Agnihotri, N.T., Streater, S., Hawkins, R.D., and Kandel, E.R. (2004). Increased attention to spatial context increases both place field stability and spatial memory. Neuron *42*, 283–295.
20. Sharp, P.E. (1999). Complimentary roles for hippocampal versus subicular/entorhinal place cells in coding place, context, and events. Hippocampus *9*, 432–443.
21. Schimanski, L.A., Peter, L., and Barnes, C.A. (2013). Tracking the course of hippocampal representations during learning: When is the map required? J. Neurosci. *33*, 3094–3106.
22. Barnes, C.A., Suster, M.S., Shen, J., and McNaughton, B.L. (1997). Multistability of cognitive maps in the hippocampus of old rats. Nature *388*, 272–275.
23. Hok, V., Chah, E., Reilly, R.B., and O'Mara, S.M. (2012). Hippocampal dynamics predict interindividual cognitive differences in rats. J. Neurosci. *32*, 3540–3551.
24. Trouche, S., Perestenko, P. V., Van De Ven, G.M., Bratley, C.T., McNamara,

C.G., Campo-Urriza, N., Black, S.L., Reijmers, L.G., and Dupret, D. (2016). Recoding a cocaine-place memory engram to a neutral engram in the hippocampus. Nat. Neurosci. *19*, 564–567.

25. Ziv, Y., Burns, L.D., Cocker, E.D., Hamel, E.O., Ghosh, K.K., Kitch, L.J., El Gamal, A., and Schnitzer, M.J. (2013). Long-term dynamics of CA1 hippocampal place codes. Nat. Neurosci. *16*, 264–266.
26. Sheintuch, L., Rubin, A., Brande-Eilat, N., Geva, N., Sadeh, N., Pinchasof, O., and Ziv, Y. (2017). Tracking the same neurons across multiple days in Ca2+ imaging data. Cell Rep. *21*, 1102–1115.
27. Meshulam, L., Gauthier, J.L., Brody, C.D., Tank, D.W., and Bialek, W. (2017). Collective behavior of place and non-place neurons in the hippocampal network. Neuron *96*, 1178–1191.
28. Rubin, A., Geva, N., Sheintuch, L., and Ziv, Y. (2015). Hippocampal ensemble dynamics timestamp events in long-term memory. Elife *4*, e12247.
29. Cai, D.J., Aharoni, D., Shuman, T., Shobe, J., Biane, J., Lou, J., Kim, I., Baumgaertel, K., Levenstain, A., Tuszynski, M., *et al.* (2016). A shared neural ensemble links distinct contextual memories encoded close in time. Nature *534*, 1–16.
30. Mankin, E.A., Sparks, F.T., Slayyeh, B., Sutherland, R.J., Leutgeb, S., and Leutgeb, J.K. (2012). Neuronal code for extended time in the hippocampus. Proc. Natl. Acad. Sci. *109*, 19462–7.
31. Mau, W., Sullivan, D.W., Kinsky, N.R., Hasselmo, M.E., Howard, M.W., and Eichenbaum, H. (2018). The Same Hippocampal CA1 Population Simultaneously Codes Temporal Information over Multiple Timescales. Curr. Biol. *28*, 1499–1508.e4.
32. Dragoi, G., and Tonegawa, S. (2013). Development of schemas revealed by prior experience and NMDA receptor knock-out. Elife *2013*, 1–24.
33. Battaglia, F.P. (2004). Local sensory cues and place cell directionality: additional evidence of prospective coding in the hippocampus. J. Neurosci. *24*, 4541–4550.
34. Bjerknes, T.L., Dagslott, N.C., Moser, E.I., and Moser, M.B. (2018). Path integration in place cells of developing rats. Proc. Natl. Acad. Sci. *115*, E1637–

E1646.
35. Yoganarasimha, D., Yu, X., and Knierim, J.J. (2006). Head direction cell representations maintain internal coherence during conflicting proximal and distal cue rotations: comparison with hippocampal place cells. J. Neurosci. *26*, 622–631.
36. Kelemen, E., and Fenton, A.A. (2010). Dynamic grouping of hippocampal neural activity during cognitive control of two spatial frames. PLoS Biol. *8*, e1000403.
37. Lee, I., Yoganarasimha, D., Rao, G., and Knierim, J.J. (2004). Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. Nature *430*, 456–459.
38. Kelemen, E., and Fenton, A.A. (2016). Coordinating different representations in the hippocampus. Neurobiol. Learn. Mem. *129*, 50–59.
39. Fenton, A.A., and Muller, R.U. (1998). Place cell discharge is extremely variable during individual passes of the rat through the firing field. Proc. Natl. Acad. Sci. *95*, 3182–3187.
40. Jackson, J., and Redish, A.D. (2007). Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. Hippocampus *17*, 1209–1229.
41. Amaral, D.G., and Witter, M.P. (1989). The three-dimensional organization of the hippocampal formation: A review of anatomical data. Neuroscience *31*, 571–591.
42. Wills, T.J., Lever, C., Cacucci, F., Burgess, N., and O'keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. Science *308*, 873–876.
43. Battaglia, F.P., and Treves, A. (1998). Attractor neural networks storing multiple space representations: A model for hippocampal place fields. Phys. Rev. E *58*, 7738–7753.
44. Tsodyks, M. (1999). Attractor network models of spatial maps in hippocampus. Hippocampus *9*, 481–489.
45. Samsonovich, A., and McNaughton, B.L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. J. Neurosci. *17*, 5900–5920.

46. Jezek, K., Henriksen, E.J., Treves, A., Moser, E.I., and Moser, M.B. (2011). Theta-paced flickering between place-cell maps in the hippocampus. Nature *478*, 246–249.
47. Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. Proc. Natl. Acad. Sci. *79*, 2554–2558.
48. Attardo, A., Fitzgerald, J.E., and Schnitzer, M.J. (2015). Impermanence of dendritic spines in live adult CA1 hippocampus. Nature *523*, 592–596.
49. Yoganarasimha, D., and Knierim, J.J. (2005). Coupling between place cells and head direction cells during relative translations and rotations of distal landmarks. Exp. Brain Res. *160*, 344–359.
50. Keinath, A.T., Julian, J.B., Epstein, R.A., and Muzzio, I.A. (2017). Environmental geometry aligns the hippocampal map during spatial reorientation. Curr. Biol. *27*, 309–317.
51. Cheng, K. (1986). A purely geometric module in the rat's spatial representation. Cognition *23*, 149–178.
52. Weiss, S., Talhami, G., Gofman-Regev, X., Rapoport, S., Eilam, D., and Derdikman, D. (2017). Consistency of spatial representations in rat entorhinal cortex predicts performance in a reorientation task. Curr. Biol. *27*, 3658–3665.
53. Julian, J.B., Keinath, A.T., Muzzio, I.A., and Epstein, R.A. (2015). Place recognition and heading retrieval are mediated by dissociable cognitive systems in mice. Proc. Natl. Acad. Sci. *112*, 6503–6508.
54. Kinsky, N.R., Sullivan, D.W., Mau, W., Hasselmo, M.E., and Eichenbaum, H.B. (2018). Hippocampal place fields maintain a coherent and flexible map across long timescales. Curr. Biol. *28*, 3578–3588.
55. Sharp, P.E., Blair, H.T., Etkin, D., and Tzanetos, D.B. (1995). Influences of vestibular and visual motion information on the spatial firing patterns of hippocampal place cells. J. Neurosci. *15*, 173–189.
56. Knierim, J.J., Kudrimoti, H.S., and McNaughton, B.L. (1995). Place cells, head direction cells, and the learning of landmark stability. J. Neurosci. *15*, 1648–1659.
57. Miao, C., Cao, Q., Ito, H.T., Yamahachi, H., Witter, M.P., Moser, M.B., and Moser, E.I. (2015). Hippocampal remapping after partial inactivation of the medial

entorhinal cortex. Neuron *88*, 590–603.

58. Kanter, B.R., Lykken, C.M., Avesar, D., Weible, A., Dickinson, J., Dunn, B., Borgesius, N.Z., Roudi, Y., and Kentros, C.G. (2017). A novel mechanism for the grid-to-place cell transformation revealed by transgenic depolarization of medial entorhinal cortex layer II. Neuron *93*, 1480–1492.
59. Kentros, C., Hargreaves, E., Hawkins, R.D., Kandel, E.R., Shapiro, M., and Muller, R. V (1998). Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. Science *280*, 2121–2126.
60. Agnihotri, N.T., Hawkins, R.D., Kandel, E.R., and Kentros, C. (2004). The long-term stability of new hippocampal place fields requires new protein synthesis. Proc. Natl. Acad. Sci. *101*, 3656–61.
61. Jeffery, K.J., Gilbert, A., Burton, S., and Strudwick, A. (2003). Preserved performance in a hippocampal-dependent spatial task despite complete place cell remapping. Hippocampus *13*, 175–189.
62. Cooper, B.G., and Mizumori, S.J. (2001). Temporary inactivation of the retrosplenial cortex causes a transient reorganization of spatial coding in the hippocampus. J. Neurosci. *21*, 3986–4001.
63. Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. Nature *532*, 459–464.
64. Morris, T., Dayoub, F., Corke, P., Wyeth, G., and Upcroft, B. (2014). Multiple map hypotheses for planning and navigating in non-stationary environments. Proc. - IEEE Int. Conf. Robot. Autom., 2765–2770.
65. Morris, T., Dayoub, F., Corke, P., and Upcroft, B. (2014). Simultaneous localization and planning on multiple map hypotheses. IEEE Int. Conf. Intell. Robot. Syst., 4531–4536.
66. Dana, H., Chen, T., Hu, A., Shields, B.C., Guo, C., Looger, L.L., Kim, D.S., and Svoboda, K. (2014). Thy1-GCaMP6 transgenic mice for neuronal population imaging in vivo. PLoS One *9*, e108697.
67. Chen, T.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., *et al.* (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. Nature *499*, 295–

300.
68. Mukamel, E.A., Nimmerjahn, A., and Schnitzer, M.J. (2009). Automated analysis of cellular signals from large-scale calcium imaging data. Neuron *63*, 747–760.
69. Zhou, P., Resendez, S.L., Rodriguez-Romaguera, J., Jimenez, J.C., Neufeld, S.Q., Stuber, G.D., Hen, R., Kheirbek, M.A., Sabatini, B.L., Kass, R.E., *et al.* (2018). Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. Elife *7*, 1–37.
70. Pnevmatikakis, E.A., Soudry, D., Gao, Y., Machado, T.A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., *et al.* (2016). Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. Neuron *89*, 299.
71. Markus, E.J., Barnes, C.A., McNaughton, B.L., Gladden, V.L., and Skaggs, W.E. (1994). Spatial information content and reliability of hippocampal CA1 neurons: Effects of visual input. Hippocampus *4*, 410–421.
72. Rubin, A., Sheintuch, L., Brande-Eilat, N., Pinchasof, O., Rechavi, Y., Geva, N., and Ziv, Y. (2019). Revealing neural correlates of behavior without behavioral measurements. Nat. Commun. *10*.