Mice can learn a cognitive map based on a stable start location and self-motion signals

Jiayun Xu[1*], Mauricio Girardi-Schappo[2,5*], Jean-Claude Béïque[1,3,4], André Longtin[1,2,3,4], Leonard Maler[1,3,4]**

[1.] Department of Cellular and Molecular Medicine, University of Ottawa, Ottawa, Ontario, Canada, K1H 8M5
[2.] Department of Physics, University of Ottawa, Ottawa, Ontario, Canada, K1N 6N5
[3.] Brain and Mind Institute, University of Ottawa, Ottawa, Ontario, Canada, K1H 8M5
[4.] Center for Neural Dynamics, University of Ottawa, Ottawa, Ontario, Canada, K1H 8M5
[5.] *Current Address*: Departamento de Física, Universidade Federal de Santa Catarina, 88040-900, Florianópolis, Santa Catarina, Brazil
***** Contributed equally to this work
****** Author for correspondence: lmaler@uottawa.ca

Abstract

Animals navigate by learning the spatial layout of their environment. Cognitive map theory asserts that they generate a metric map of their environment that permits short cuts, but tests for a cognitive map are inconclusive. We investigated spatial learning of mice in an open maze where food was randomly hidden in one of a hundred holes. Mice leaving from a stable entrance learned to efficiently navigate to the food without the need for distal landmarks. We successively trained the mice to first find food at Site A then B. On a probe trial, the mice checked for food at B and then took a novel short-cut to A. Detailed analyses of mice trajectories, active hole checks and varying the entrance site demonstrated that the mice had formed a cognitive map by using the stable entrance location to anchor an egocentric map based on path integration of self-motion signals.

Introduction

Animals must learn the spatial layout of their environment to successfully forage and return home. Navigation can be based on moving towards a beacon[1] or learning a series of motor actions in response to proximal cues [2]. Experimental studies have also used one[3] or many[4] distal landmarks (allothetic cues) to provide an allocentric coordinate frame for learning trajectories to a hidden location. A critical test of whether an animal is flexibly using an array of distal landmarks for navigation is whether, having learned to reach a target from one start point, it can generalize and reach the target from a different starting location. Eichenbaum et al[4] used this protocol (water maze) and found that rats successfully navigated from a novel start location to a hidden platform, but Benhamou [5] found that rats failed the novel start location test.

A second source of spatial information comes from idiothetic cues and provides a moving egocentric reference frame via path integration of self-motion signals[6]. Behavioral and physiological data shows that the egocentric frame can be anchored to the external coordinates[7]. The combination of allothetic and idiothetic inputs is hypothesized to generate a cognitive map[7] that provides a coordinate frame permitting trajectories between arbitrary locations to be computed. A test for a cognitive map is an animal's ability to take short cuts by using the egocentric or allocentric coordinates of remembered sites.

We developed a circular open maze where a mouse had to find food randomly hidden in one of 100 holes. Large distal cues were visible when the mice left their home cage but, when the mouse's starting point was changed from trial to trial, it could not learn the food location. When trained on one start site, the mice learned to reach the food location via efficient trajectories and did not require the cues. Mice failed the different start site test - their trajectories were unchanged after leaving from a different entrance and no longer led to the food site. These results suggest that mice had not developed a cue-based allocentric map but were relying only on egocentric cues.

To test for a cognitive map, we developed a protocol to determine whether mice trained from a static start site could compute a short cut between two memorized food sites. We trained the mice to find food at one site (A) and, subsequently, at a second site (B). In a probe trial (no food), the mice first ran to B and then took a novel route from B to A, i.e., they took a short cut. We conclude that mice, in our open maze, will assume that they are leaving from a stable start location and learn a cognitive map using only self-motion cues.

Results

We designed a behavioural framework, the Hidden Food Maze (HFM), to tease apart the roles of idiothetic and allothetic cues for navigation in a foraging task (Fig. 1a-d). Mice are trained to find a food reward hidden inside one of 100 holes in a large, circular open maze. The HFM has several unique features such as rotational symmetry to eliminate geometric cues and control for distal cues, movable home cage entranceways, and experimenter-controlled visual cues (see Methods). These features allow us to control whether mice orient to allothetic or idiothetic cues, enabling us to investigate the respective role of landmarks versus path integration in spatial navigation. We start with visual cues and analyze the effect of randomizing the entrance taken by the mouse on successive days (Fig. 1e), versus maintaining a static entrance throughout training (Fig. 1f). We analyze the trajectories in quantitative detail, including the hole checks used to locate the food. Finally, we investigate the effect of switching the food site on the trajectories and hole checking (Fig. 1g). We developed a method to infer the mouse's target estimation

vector (TEV) of the food location. This is used to demonstrate the learning of target direction and the use of short cuts in the two-targets condition (see Methods).

Mice use idiothetic sensory input but do not utilize 2D wall cues for spatial learning

1. Random Entrance experiments.

Mice were trained with their entrance changing randomly from trial to trial, so their only reliable orienting cues were the wall cues. The configuration of the 4 cues varied relative to the entrance, but the food was always placed in the same hole and nearest the X-shaped landmark (Fig. 2, Methods). Individual mice showed persistent complex search trajectories throughout the learning period (Figure 2a, b). The latency to reach the food target was significantly negatively correlated with trial number (statistics in Fig. 2d) but with a non-significant difference between the first and last trials (N = 8; First trial mean/SD: 313.92 s/566.74 s, Last trial mean/SD: 34.88 s/29.14 s; P = 0.21582). Despite the decrease in latency, mice did not appear to exhibit spatial search strategies (Fig. 2c) given that the time spent in each quadrant during the probe trial is similar (target quadrant = 28.78%, other quadrants = 24.93%, 28.04%, 18.25%). This suggests latency decreased for reasons independent of spatial learning.

We found a corresponding significant positive correlation of speed and trial number (Fig. 2d, e) with strong differences between first and last trials (N = 8; First trial mean = 8.12cm/s, SD = 4.28, Last trial mean = 18.50cm/s, SD = 6.06; P = 0.00018). There is no significant correlation of normalized trajectory length (measured length/direct length to food) and trial number (Fig. 2f). The first and last trajectory lengths were not significantly different either (N = 8, First Trial mean = 1175.95, SD = 1048.99, Last Trial mean = 500.94cm, SD = 252.81; P = 0.12604). We hypothesize that the mice learned to search more rapidly, and reducing the home to food time, consistent with the random trajectories before and after learning (Fig. 2a vs Fig. 2b; Extended Data Figs. 1,2).

We next examined hole checking as an active sensing indicator of the mouse's expectation of the food location (Methods Fig. 2; Supplementary Information Hole check Video). We expected that the mice would increase the number of hole checks near the expected food location, or along their path toward it. However, the hole checking density (number of checks/distance travelled) remains constant between the first and second halves of any trial, suggesting mice did not anticipate where the food is placed (Fig. 2g). Mice also did not increase their hole sampling rate as they approached the target, again suggesting no expectation of its location (Fig. 2h, i). Mice do not appear to learn a landmark-based allocentric spatial map.

2. Static Entrance Experiments.

We tested whether the mice might learn an allocentric spatial map if they had a single stable view of the landmarks and could associate one cue configuration with the food location (Fig. 3, see Methods). With static entrances, mice greatly improve their trajectory efficiency after learning (Fig. 3a, b). During early learning, mice show random search trajectories that transform to stereotyped trajectories towards food by late learning (Fig. 3b). Typical trajectories are (1) directly to the food, (2) initially displaced towards the maze center then returning to a food-oriented trajectory, or (3) initially along the wall and then turning and heading to the food (Fig. 3b). The normalized distances were now close to 1 (Fig. 3f).

During the probe trial, mice spent the most time searching in the target quadrant (target quadrant = 44.26%, other quadrants = 19.38%, 16.55%, 19.81%), demonstrating acquisition of spatial knowledge (Fig. 3c). The latency to target was strongly negatively correlated with trial number (Fig. 3d) and the last trial latency was significantly reduced compared to the first trial (N = 8; First Trial mean = 224.24s, SD = 170.08, Last trial mean = 6.67s, SD = 3.83; P = 0.011649). The speed was significantly positively correlated with trial number (Fig. 3e) and the ending speed significantly greater than the initial speed (N = 8, First trial mean = 9.45 cm/s, SD = 3.95, Last trial mean = 23.10cm/s, SD = 5.39; P = 0.000099). The normalized trajectory length was now strongly negatively correlated with trial number (Fig. 3f) and the final length greatly reduced (N = 8; First trial mean = 1893.51cm, SD = 1539.15, Last trial mean = 140.59cm, SD = 59.73; P = 0.020332; Extended Data Figs.1,2).

After learning, hole checking density varies across trajectories and differs from the early learning trials. During individual trajectories, hole checking significantly decreases with trial number in the first half of trajectories (Fig. 3g). The average distance of the checked holes to the food decreases significantly for the second half of the trajectories (Fig. 3h). The continued decrease results in hole checks happening much closer to the food when learning is complete (Trial 14) since $d_{(checks \rightarrow food)}$ is almost zero. The most substantial effect of learning is that the probability of hole checking increases as the mouse approaches the food (Fig. 3i). These results (Fig. 3h,i) suggest that the mice may be predicting the food location and checking their prediction as they near food.

Both random and static entrance experiments start with random trajectories, long travel times and trajectory lengths and high numbers of checked holes. These measures remain nearly constant throughout the random entrance trials and covary within a large range (Extended data Fig. 1). In the static entrance case, these features decrease together after learning (Extended data Fig. 2). Comparing geometric features of their trajectories, the mice in the static entrance

experiment significantly optimize their trajectory towards the food with their travel distance approaching the straight line between the entrance and food (Extended Data Figure 1d, e, f).

These results demonstrate that mice can learn the spatial location of the food target from a static entrance. We checked whether the performance gain can be attributed to the cues by manipulating the mouse's entrance. Well-trained static entrance mice were subject to rotation of their home cage by 180°, +90°, or -90°. If the mice were using the distal landmarks, we would expect them to still find the correct food location. If they relied on idiothetic cues, they should travel to the rotationally equivalent location (REL). The REL is defined as the location where the target would be if its position were invariant to the entrance used in a trial. In other words, going to the REL means that a trajectory is anchored to the mouse's home cage and not to the entrance it was using during training. As illustrated in Fig. 1f, the mouse's trajectory rotates 180° with its home cage and points to the REL.

Following all rotations, mice searched at the REL instead of the correct location for the food reward - they had not used landmarks to compensate for rotation (Fig. 3j, k, l). Mice were quick to reach the REL targets (N = 7, one mouse was excluded because it did not go precisely to the REL; Mean Latency = 8.78s, SD = 2.27). There is no significant difference between latency to reach the target during late training (N = 7, Mean = 6.94s, SD = 2.27) and the REL location after rotation (P = 0.43). The mice failed the novel start location test and are not using allothetic cues for navigation.

3. Controls

A) It is possible that mice, like gerbils[3], choose one landmark and learn the food location relative to that landmark (e.g., the "X" in Fig. 3b); in this case, the mice would only have to associate their initial orientation to that landmark without identifying the image on it. We used two additional controls to assess this possibility. Mice were trained without any cues on the wall, eliminating the possibility that they provided orienting guides. The learning curve of mice with or without the 2D wall cues were not significantly different (Extended Data Fig. 3a, b).

B) Mice, well trained in the lighted maze, were given trials in complete darkness. These mice still showed the same learning curve compared to untrained mice or mice following rotation (Extended Fig. 3c, d).

C) Finally, to completely rule out both landmark and possible olfactory cues, we trained and tested mice in total darkness. Head direction tuning may be impaired in sighted mice navigating in darkness[8], but spatial learning was not affected in our experiments. Mice can use localized odor sources as landmarks for spatial learning [9], and might therefore use odor gradients from

their home cage as a cue. We eliminated odor cues by applying negative pressure inside the home cage via fans that drew in outside air from the cage bottom and from the maze and vented cage air out the top (see Methods). The mice were still able to learn the food location in the absence of visual and olfactory cues (Extended Data Fig. 3e,f). These experiments demonstrate that the mice learn to efficiently navigate from a fixed start location to the food location using only idiothetic cues and path integration.

Learning trajectories and active sensing

A comparison of trajectories in the random and static entrances conditions was performed using "step maps" across all mice in one trial of a learning protocol (see Methods: Step Map Calculation and Extended data Fig. 4). Step maps discretize the maze into a lattice and compute the probability of a mouse moving one step from its current lattice location in a forward/backward or left/right direction. This results in a spatial gradient of the probabilities of moving from a lattice location to one of its adjacent sites. The Target Estimation Vector (TEV) is the vector sum of all steps in the map, giving an estimation of where the net trajectory is pointing for all the mice combined in each trial. The TEV quantifies the evolution of the direction of trajectories during learning.

Random entrance training. The mice's step distributions before and after training are similar with the step map pointing randomly at every lattice site. The TEV point in a random direction, confirming that, after training, mice homogenously search the maze (Fig. 4a, b; Extended Data Fig. 4a,b). Mice exhibited a uniform hole checking pattern in early trials (Fig. 4c) with many checks near the wall. During late learning, mice preferred to search in the location where food last appeared from their previous entrance (Fig. 4d).

Static entrance training. In the static entrance experiment multiple step trajectories become biased towards the food; trajectories may initially deviate from the food vector but tend to return to it (Fig.4f, Extended Data Fig. 4d). The angle between the TEV and the food vector diminishes during training and TEV eventually points almost directly to the food lattice site (Fig. 4f; Extended Data Fig.4d,e – the error is 11 cm, which is the size of a site in the lattice). Mice developed a strong preference for searching along the TEV and at the correct target location (Fig. 4g,h; also see Fig. 3i).

Mice can use a short cut to navigate between remembered targets

Our results establish that mice learn the spatial location of a food reward using idiothetic cues and path integration. Have the mice learned an egocentric spatial map, or a more flexible cognitive map? A test of cognitive mapping would be if mice can take a short cut and navigate

from one remembered location to another along a novel, unreinforced path. The “two food location experiment” set out to test this possibility by training mice to travel to food from their home cage to two very differently located sites. The two locations are trained sequentially so that one site is trained first followed by training on the second site. Only one hole is filled with food in any given training trial. Mice are then tested to see whether they can travel between the two remembered locations during a probe trial (food is absent) despite no prior training or unreinforced experience on the short cut route (Fig. 1g).

Mice successfully learned the location of Target A, evidenced by their direct search trajectories and significant decreases in distance traveled (Fig. 5a), with the same performance as in the static entrance case. During the learning of Target A, the distance of the mouse to the future location of target B was always larger than that expected by chance (Extended Data Fig. 5). In Fig. 5a, the mouse passes within 16 cm of the future Target B site in an early learning trial, but this closest approach does not constitute a B->A trajectory. The first probe trial has no food in Target A, but the mouse initially still heads for Target A and, finding it empty, re-initiates a random search for food (Fig. 5b,c).

After the probe trial, mice are trained to learn Target B. During this training, 3 mice (Mouse 33, 35, 36) first learned a route to Target B but, on subsequent searches, they would sometimes first check Target A and then travel to Target B along a direct short cut route (see Discussion).

After learning the location of Target B (Fig. 5c), food was then omitted from this target as well for a second probe trial (labeled Probe B->A), meant to check whether the mice would take shortcut routes between B and A. As illustrated in Fig. 5d and fully illustrated in Extended Data Fig. 6, 5 of 8 mice were observed to go from Target B to Target A via varying trajectories. Two of the remaining three mice went first to Target A and subsequently to Target B via direct or indirect routes. These mice had previously taken direct routes to Target B, showing that going to Target A first is not due to a learning deficiency. The remaining mouse (mouse 34) went from Target B to the start location and, from there, to Target A. This mouse previously had taken the B -> A route during training. In all cases, the mice clearly remembered the location of the newly learned Target B and of Target A, even though the latter was taken and rewarded 4 days previously.

To compare the distance traveled towards target locations of different lengths, we normalized the total travel distance from the home cage to each of the targets and between targets, i.e., divided the length of the path taken by the actual direct distance (Fig. 5e). The normalized B->A and A->B trajectories are comparable in directness (i.e. have similar normalized distance) to those taken by well-trained mice travelling towards targets A and B from the home cage start site. Notably, the normalized short cut trajectory is also shorter than the trajectories of the mice

traveling between two arbitrary points in the same trial (Fig. 5e). The heading angle towards the food and distance of the mouse to the food line are also decreased while the mice are travelling between targets B and A compared to prior to learning (Fig. 5f,g).

We were concerned that the mice, when learning target B, would take routes that passed near target A and went to target B. In this case, the reward at target B might result in learning a "near A to B" trajectory. We therefore calculated, for all target B training trials, the mouse's closest distance to a random point and its minimum distance to the previous target A. As seen in Extended data Fig. 5, the mice do not systematically pass near A on route to B, since the minimum distance to A is greater than what is expected by chance. We did the same calculations for learning target A, but now considering the minimum distance to the future location of the B target. In this case the mice also did not pass closer to B than expected by chance. The trajectories stayed, on average, approximately 40 cm away from target B. In other words, the mice did not learn, by chance, direct unreinforced routes from "near B to A".

The step probabilities again reflect the mice leaving and returning first to Target A TEV (Fig. 6a; Extended Data Fig. 7a) then to Target B TEV (Fig. 6d, Extended Data Fig. 7c). On the last probe trial, they further reflect leaving and returning to the TEV directed from Target B towards Target A (Fig. 6g; Extended Data Fig. 7d,e).

Hole checking analysis revealed that trained mice first check for holes in the vicinity of target A as expected (Fig. 6b), and this continues on the first trial with food at target B (Fig. 6e). Hole checking then correctly shifts targets when the mice have learned the target B location (Fig. 6f). During the Probe B→A trial, hole checking clusters around both target A and B as well as the TEV connecting B and A (Fig. 6h).

<u>Control.</u> As a stringent control, we first trained the mice on the two-target protocol as described above. Before the final probe trial we rotated the mice by 180° as described for static entrance experiments (Fig. 3j-k). The mice trajectories now first went to the REL for Target B followed by going directly to the REL for Target A (Extended Data Fig. 8). The mice acted in a way consistent with independent responses in the two food site probe trials and the entrance rotation experiments, including the pattern of hole checking. Mice can indeed take a novel short cut and have therefore learned a cognitive map based on a fixed starting point and egocentric input.

Discussion

We have shown that mice are unable to learn the location of food when their entrance is randomly assigned on successive trials. When trained with a static entrance, mice rapidly learn the food location and would appear to need consistent association of landmark and food location

for learning. The same results for static entrances are obtained when distal cues are omitted, or the mice are trained and/or tested in the dark and in the absence of olfactory cues. The learned trajectory lengths are also consistent and near optimal in all these cases. We conclude that the mice do not rely significantly on landmarks for our spatial learning task.

Our detailed analysis of the local trajectory probabilities demonstrates that the mice have learned to estimate the Start to food trajectory (food vector, Extended Data Fig. 1) and that departures from this vector are compensated for, presumably by path integration signals. We considered that the mice might be using unidentified cue(s) to guide their accurate trajectories. We therefore, in a final probe trial, rotated the mouse's home by 180°, 90° or -90°. The mouse ran to the relatively equivalent location (REL), i.e., to the spot where the food was estimated to be relative to the home position had the mouse not been rotated (Fig. 3j-l). This result was reproduced even within the two-targets protocol. We conclude that (a) the mice have an accurate estimate of the food location when they leave home, (b) they continuously update and correct their location via a path integration mechanism and (c) the resulting egocentric spatial map is anchored to their home cage location.

Rodents make active sensing motions [10] and head-scanning has been directly linked to the formation or strengthening of hippocampal place fields [11]. We rarely observed active sensing motions in our experiments. The mice did actively check holes and, after learning in the static condition, the probability of hole checking increased as the mouse neared the food site (Fig. 3h,i), reinforcing our conclusion that the mice's trajectories are based on an initial estimate of food location. We interpret the increase in hole checking as evidence for accumulation of error in the path integration mechanism [6]. Active sensing may be more generally connected to spatial learning. Work on electric fish has demonstrated that electrosensing scanning motions are used to identify landmarks and thereby anchor egocentric signals that efficiently guide the fish to the food site via path integration [12]. Like the mouse, the fish's sampling increases as they approach the learned food site, presumably due to error accumulation during path integration[13].

In the critical two food sites experiments, mice successively learned to find food first at site A, then at site B, and we tested their behavior when food was missing from both sites. Five of eight mice took a short cut trajectory from site B to the earlier (4 days) learned site A (Fig. 5a-d). The normalized distance travelled in the short cut trajectories and all other measured variables were like those measured for simple learning to site A or site B (Fig. 5e-h). The same results applied when the mice's home was rotated for the final probe trial (Extended Data Fig. 8), i.e., the mice do not use any allothetic cues when making the short cut. We were also concerned that the short cut was not novel but learned after "unrewarded exposure to the shortcut route"[14]. However,

our calculations showed that the mice did not by chance learn direct routes between sites A and B (Extended Data Fig. 6), i.e., the B->A shortcut was novel. Bennett[15] notes that path integration "allows an animal to make a straight-line return to the starting point" and this is not evidence for a cognitive map. In our experiments, the novel short cut was not to home but between other previously memorized locations. We conclude that the mice can learn a cognitive map in a manner consistent with previous hypotheses[7, 16] – they combine knowledge of a stable location with idiothetic signals to generate an egocentric-based map that efficiently connects their start point with a food site and connects previously learned food sites.

Previous studies showed that mice can use distal cues under different experimental conditions, such as in a one-dimensional T-maze, linear track, or when the cues consist of multiple highly salient cues [17-20]. Their failure to use such cues in our experiments may be due to the lesser saliency of our visual cues. Alternatively, it may be due to our pre-training procedure. Mice were hesitant to explore the center of the maze and, to expedite training, we omitted landmarks and placed treats in numerous holes (see Methods). We propose that mice learned to ignore the landmarks and initially employed a random search strategy to find food.

Rodents can learn many ways to navigate to a reward or home. Mice use local cues to learn the structure of a complex labyrinth to optimally find a water reward[2]. Mice, in a partially open maze, can learn multi-step routes using subgoals to return home or find food[21]. We have shown that mice in an open maze can combine the location of a stable home and self-motion cues to generate a cognitive map. The same options are available to humans. Zhu et al[22] showed that subjects used active sensing eye movements to learn to efficiently navigate to hidden goals in a virtual maze, a case of strict allocentric spatial learning. Landau et al[23, 24] described a blind toddler that took a short cut between two objects previously reached from its mother – path integration equivalent to our two food site experiments. The mice short cuts formally amount to vector subtraction: $\overrightarrow{B \rightarrow A} = \overrightarrow{Home \rightarrow A} - \overrightarrow{Home \rightarrow B}$, something that is precisely captured by our TEV calculation. A recent theoretical analysis concluded that a recurrent network trained to perform path integration can develop a grid cell-like representation of spatial relations. The network was then capable of vector-based navigation including short cut behaviors[25]. We conclude that the cognitive map behavior we observe is within the capabilities of the circuitry and spatial responsiveness of neurons within the entorhinal and hippocampal cortices.

Funding

This work was supported by the Canadian Institutes for Health Research Grant # 153143 to AL and LM, a Brockhouse award (493076-2017) to AL, LM and J-C B, an NSERC award (RGPIN/06204-

2014) to AL, an NSERC award to LM (RGPIN/2017-147489-2017) and a grant from the Krembil Foundation to AL, LM and J-C B.

Acknowledgements

We would like to thank Dr. Érik Harvey-Girard for technical support.

Supplementary Video

The supplementary video can be downloaded from:
https://www.dropbox.com/s/ikts0dn1qcog4oi/mouse_36_Probe2_hole_check_synchronized.mp4?dl=0

The video illustrates a two-food site experiment including hole checks. The mouse was first trained with food at Site A (6 days) and, after training was complete, trained with food at Site B (3 days). The video was taken on a probe trial (Day 10) with no food at either Site A or Site B. The mouse is observed to proceed along a TEV to Site B with a departure from the TEV at the second hole check, followed by a third hole check as it loops back to Site B. After hole checking Site B, the mouse goes on a TEV to Site A with a departure after its third hole check. It then returns to the TEV with further hole checks. Just before reaching Site A the mouse turns and moves very slowly to a hole 16 cm from Site A; this hole check was missed by our algorithm and had to be manually added for statistics. The mouse spends a great deal of time at this hole, in spite of the fact that it had never contained food. We hypothesize that this was the hole the mouse's cognitive map predicted to be Site A. The mouse subsequently continued to check Site A and a site near the wall. At this point the experiment was terminated. Our "wash and rotate" protocol to eliminate any odor trails implies that floor of the maze would have been washed and rotated 108 times from the first Site A learning trial to the final Probe trial when Site B and Site A were empty.

Data availability

All codes and data are available in https://github.com/neuro-physics/mouse-cogmap

Figures

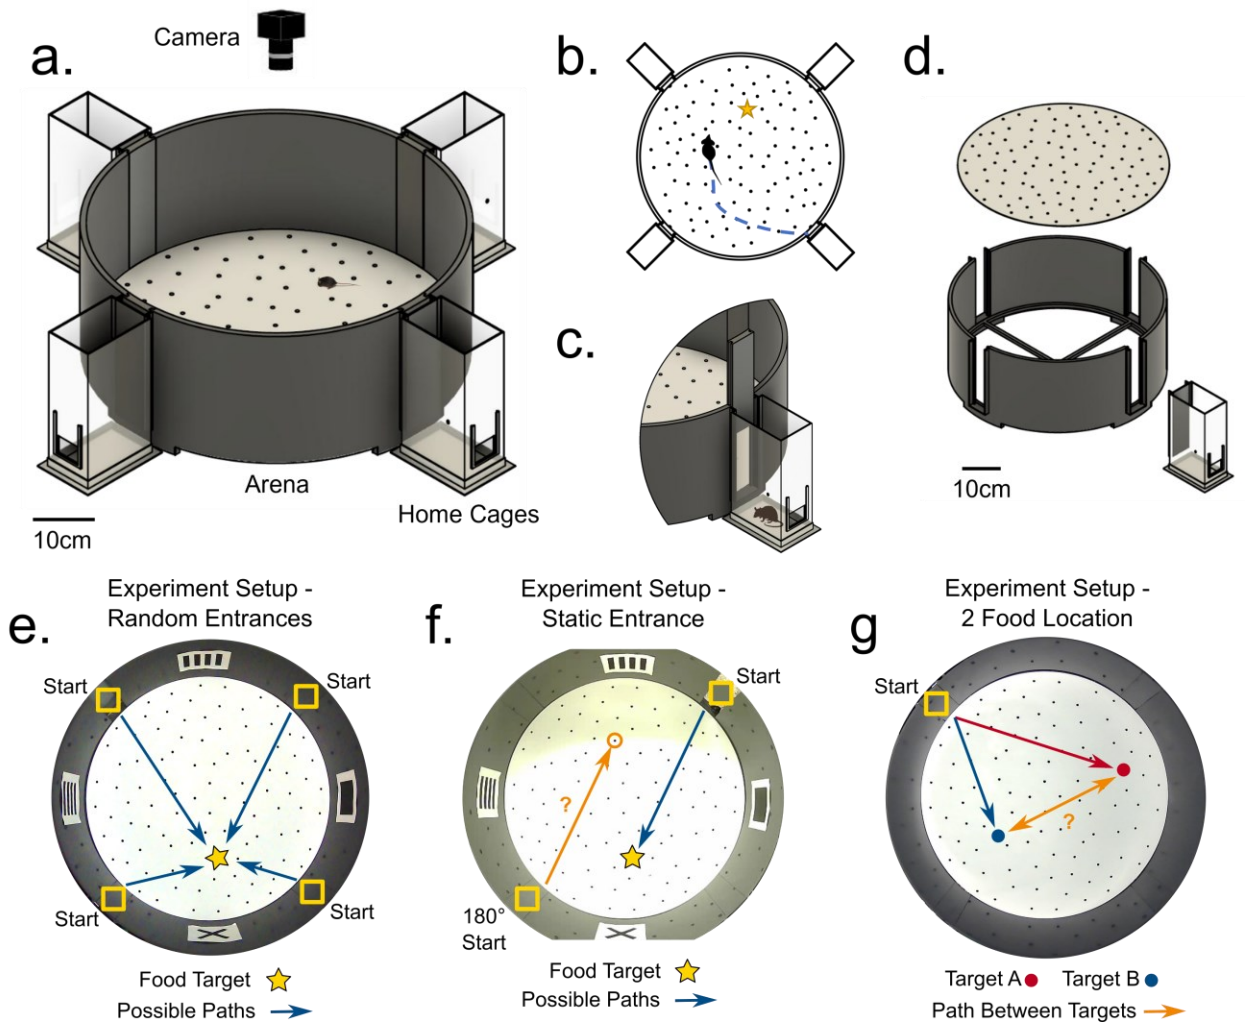

Figure 1. **Schematic of the Hidden Food Maze and the experimental setups. a.** Schematic of the Hidden Food Maze. The floor of the arena is 120cm in diameter, and the walls are 45cm tall. The home cage has a 10.5 x 6.5 cm floor area. **b.** Schematic view (top-down) of a mouse searching for hidden food (star symbol) after exiting the home cage. **c.** Detailed view of how the door slides upward, allowing mice to enter the arena without handling. **d.** Diagram of how the floor and the home cages of the arena can be separated. The subfloor containing mouse food is not illustrated. **e**. Diagram of the setup of the "random entrances" experiment. Mice can enter randomly from any of the four entrances (blue arrows) to search for a food reward (star). **f.** Diagram of the setup

of the “static entrances” experiment. Mice enter from a fixed entrance (blue arrow) to search for a food reward (star). Following a 180°, rotation of the home cage, a hypothesized path is shown if the mouse learned the food location using path integration alone (orange arrow) and ignored the landmarks. The mouse would go to the Rotationally Equivalent Location (REL, orange circle). **g.** Diagram of setup of the “two food location” experiment. Mice enter from a fixed entrance to search for a food reward (red arrow and Target A). Mice then are trained to find food in a different location (blue arrow and Target B). Possible direct paths between the two targets are evaluated during probe trials (orange arrow).

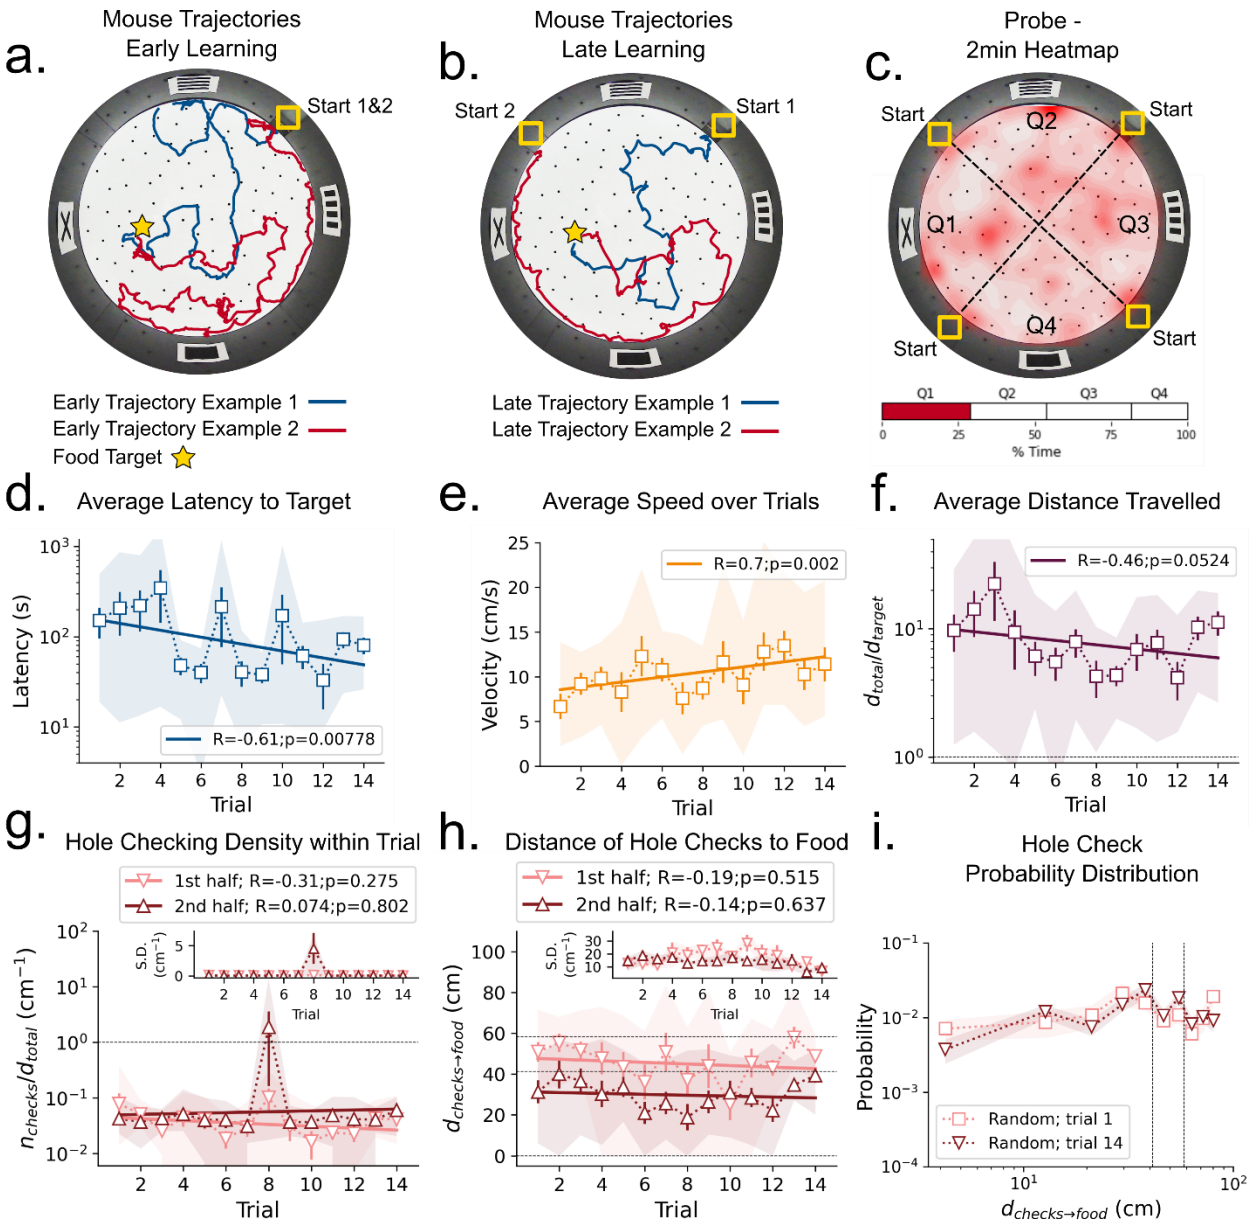

Figure 2. Mouse Spatial Learning with Changing Entrances. a. Two examples of mouse search trajectories during early learning when the entrance changes from trial to trial. The trajectories from one entrance are irregular and vary unpredictably across trials. Trial = 3. Star = target location. Yellow Square = entrance site **b.** Two examples of mouse search trajectories during late learning after leaving from different entrances. Note the similar variability of trajectories after many trials. Trial = 14, Star = target location. Yellow squares = entrance sites **c. (Top)** Heatmap of the first 2min of a probe trial done after trial 18. **(Bottom)** Stacked Bar chart showing percentage of time spent within a 15cm radius of the target (red) and its rotationally equivalent position in the other three quadrants (white). N = 8. **d.** Average latency to reach target. A significant reduction in latency is seen ($p<0.05$). Plotted with logarithmic y-axis scale. N = 8. Error bars = SE.

Shaded area = data range. **e.** Average speed during trials. The corresponding significant increase in speed is seen. N = 8. Error bars = SE. Shaded area = data range. **f.** Average normalized distance mice traveled to reach target, plotted with logarithmic y-axis scale does not significantly reduce over trials ($p > 0.05$). N = 8. Error bars = SE. Shaded area = data range. **g.** Hole-checking density (number of hole checks per distance traveled) in each half of the trajectory. The density remains constant for both halves and across trials, suggesting that mice remained uncertain as to the food location. **(inset).** The standard deviation (S.D.) of the density over the mice sample remains constant for both halves. N = 8. **h.** The average distance of the checked holes to the food $d_{(checks \rightarrow food)}$ remains almost constant across trials. Horizontal lines are just guides to the eye. **i.** The probability density profile of checking a hole versus the distance to the food $d_{(checks \rightarrow food)}$ for the first and last learning trials (the corresponding averages over trials are in panel h). The distribution remains unaltered. Vertical dotted lines mark the same distances from the horizontal lines in panel h.

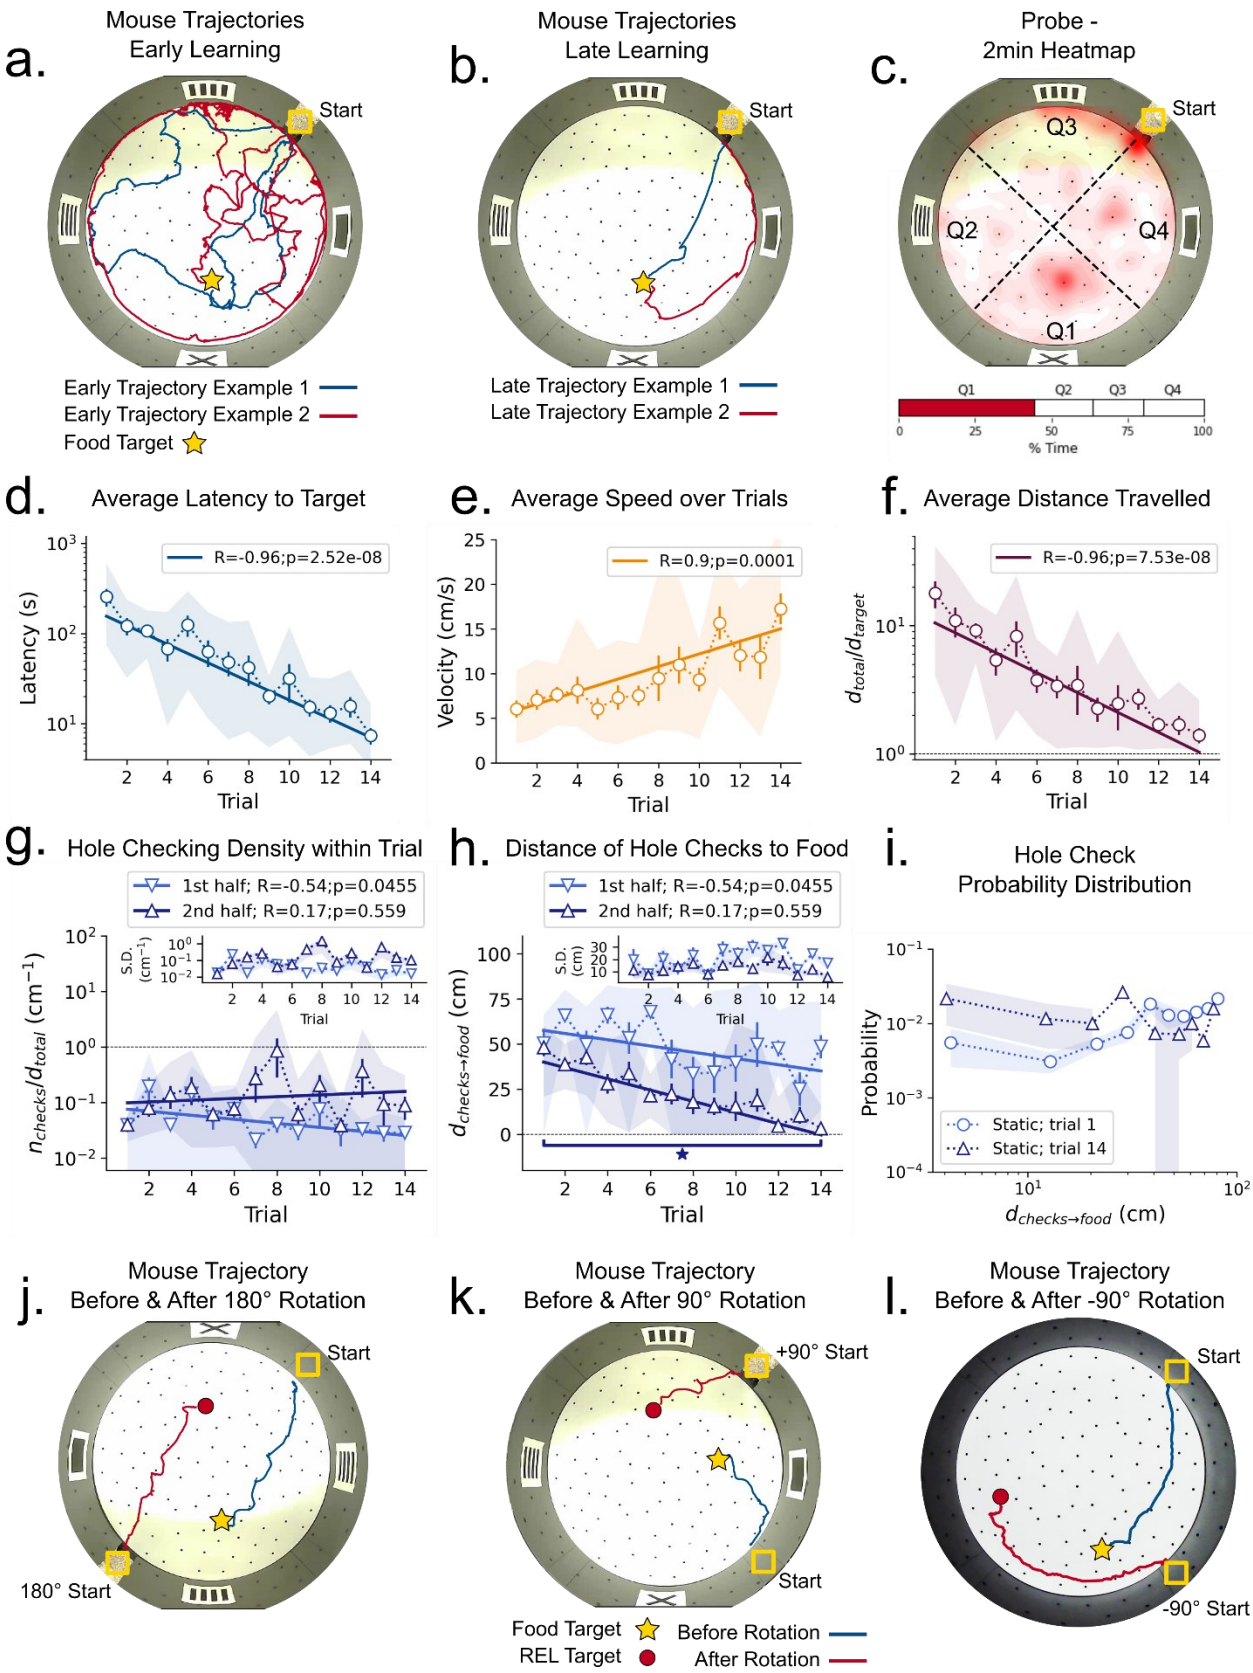
a.
Mouse Trajectories
Early Learning
Start
Early Trajectory Example 1
Early Trajectory Example 2
Food Target
b.
Mouse Trajectories
Late Learning
Start
Late Trajectory Example 1
Late Trajectory Example 2
c.
Probe -
2min Heatmap
Start
Q1
Q2
Q3
Q4
% Time
d.
Average Latency to Target
R=-0.96;p=2.52e-08
Latency (s)
Trial
e.
Average Speed over Trials
R=0.9;p=0.0001
Velocity (cm/s)
Trial
f.
Average Distance Travelled
R=-0.96;p=7.53e-08
Trial
g.
Hole Checking Density within Trial
1st half; R=-0.54;p=0.0455
2nd half; R=0.17;p=0.559
S.D. (cm⁻¹)
Trial
h.
Distance of Hole Checks to Food
1st half; R=-0.54;p=0.0455
2nd half; R=0.17;p=0.559
S.D. (cm)
Trial
i.
Hole Check
Probability Distribution
Probability
Static; trial 1
Static; trial 14
j.
Mouse Trajectory
Before & After 180° Rotation
Start
180° Start
k.
Mouse Trajectory
Before & After 90° Rotation
+90° Start
Start
Food Target
REL Target
Before Rotation
After Rotation
l.
Mouse Trajectory
Before & After -90° Rotation
Start
-90° Start

Figure 3. Mouse spatial learning with static entrances. a. Two examples of mouse search trajectories during early learning. Random and highly variable trajectories similar to those seen in the Random Entrance experiments. Trial = 3, Star = target location. **b.** Two examples of mouse search trajectories during late learning. Trajectories are clearly directed towards the food, although there is still variation of the trajectories across trials; trajectories are typically directly towards the food or first along a wall and turning towards the food. Trial = 14, Star = target location. **c. (Top)** Heatmap of the first 2min of a probe trial done after trial 14. **(Bottom)** Stacked bar chart showing percentage of time spent within a 15cm radius of the target (red) and its rotationally equivalent position in the other three quadrants (white). N = 8. **d.** Average latency to reach target. N = 8. Error bars = SE. Shaded area = data range. **e.** Average speed during trials. N = 8. Error bars = SE. Shaded area = data range. **f** Average distance mice traveled to reach target, normalized. N = 8. Error bars = SE. Shaded area = data range. **g.** Hole-checking density (number of hole checks per distance traveled) in each half of the trajectory; larger means more uncertainty. A significant decrease across trials was found in the first half of the trajectory ($p<0.05$), and an increasing trend in the second half. **Inset:** The standard deviation (S.D.) over the mice sample of the density is larger in the second half. **h.** The average distance of the checked holes to the food $d_{(checks \rightarrow food)}$ decreases across trials for both halves of the trajectory. After learning (trial 14), the hole checks happen much closer to the food since $d_{(checks \rightarrow food)}$ is almost zero, even though there are more hole checks per distance traveled. **i.** The probability density profile of checking a hole versus the distance to the food $d_{(checks \rightarrow food)}$ for the first and last learning trials (the corresponding averages over trials are in panel h). After learning (trial 14), the probability of hole checks is larger closer to the food – a feature that does not appear in the Random Entrance experiments. **j.** Mouse search trajectory before (blue) and after (red) a 180° rotation after training. Mice visit the REL target, indicating they are using path-integration. **k.** Example of a mouse search trajectory before and after a 90° rotation. **l.** Example of a different individual mouse search trajectory before and after a -90 rotation. Note that, in this case, we removed the landmarks. Nevertheless, the mouse still navigated directly to the REL.

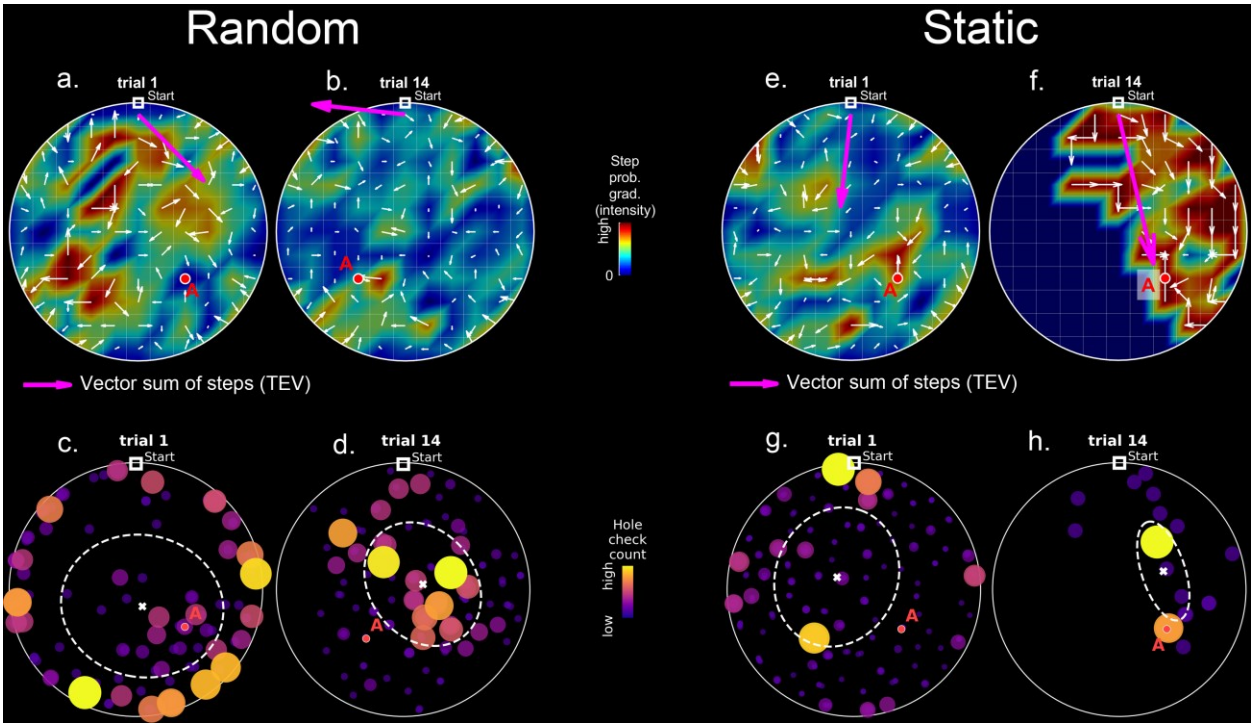

Figure 4. Step maps (top) and active sensing (bottom) for random and static entrance. Each panel on the top row corresponds to the one immediately below it. **Step maps (top row)**: the color and arrows indicate the most probable route taken from a particular location (red means more probable; blue indicates equal probability of heading to any direction) – see Step Map Calculation in Methods, and Extended data Figure 4 for more details in the Extended Data. **Active sensing (bottom row)**: the size and color of each circle indicates the normalized frequency that the hole in that location was checked (larger and lighter yellow indicates higher frequency); the ellipse centered in x marks the covariance between the x and y coordinates of the hole-checking events. In the Random Start cases, the food was always at the same location, but the mouse started at different entrances. They appear to be at different locations (b&d versus a&c) because the images were aligned, for ease of viewing, so as to appear to start at the same entrance. **a and b.** Step maps indicating the most used trajectories in the random entrance case. Both trials 1 and 14 illustrate homogenous trajectories that do not change with learning. The Target Estimation vector (TEV, purple arrow) points in a random direction and not the food site even in the late trials. **c and d.** Active sensing profile evolution over the learning trials for the random entrance experiment aligned according to mouse perspective (see Experiment Alignment in Methods and Extended data Figure 4). Hole-checking events accumulate near the center of the arena in late learning trajectories, but they are not centered over the target marked with a red (A); rather appear to be centered over where the food was located after a previous search from a different entrance. **e and f.** Step maps for the static entrance case. After learning, there are well-defined

flows for the trajectories. The TEV points directly to the food site (red A label) and most probable trajectories may deviate from the TEV (in either direction) but then return to it. **g and h.** Active sensing profile evolution in the static entrance experiment. After learning, hole-checking accumulates on and around the food line and is greatest near the food site (see Fig. 3i).

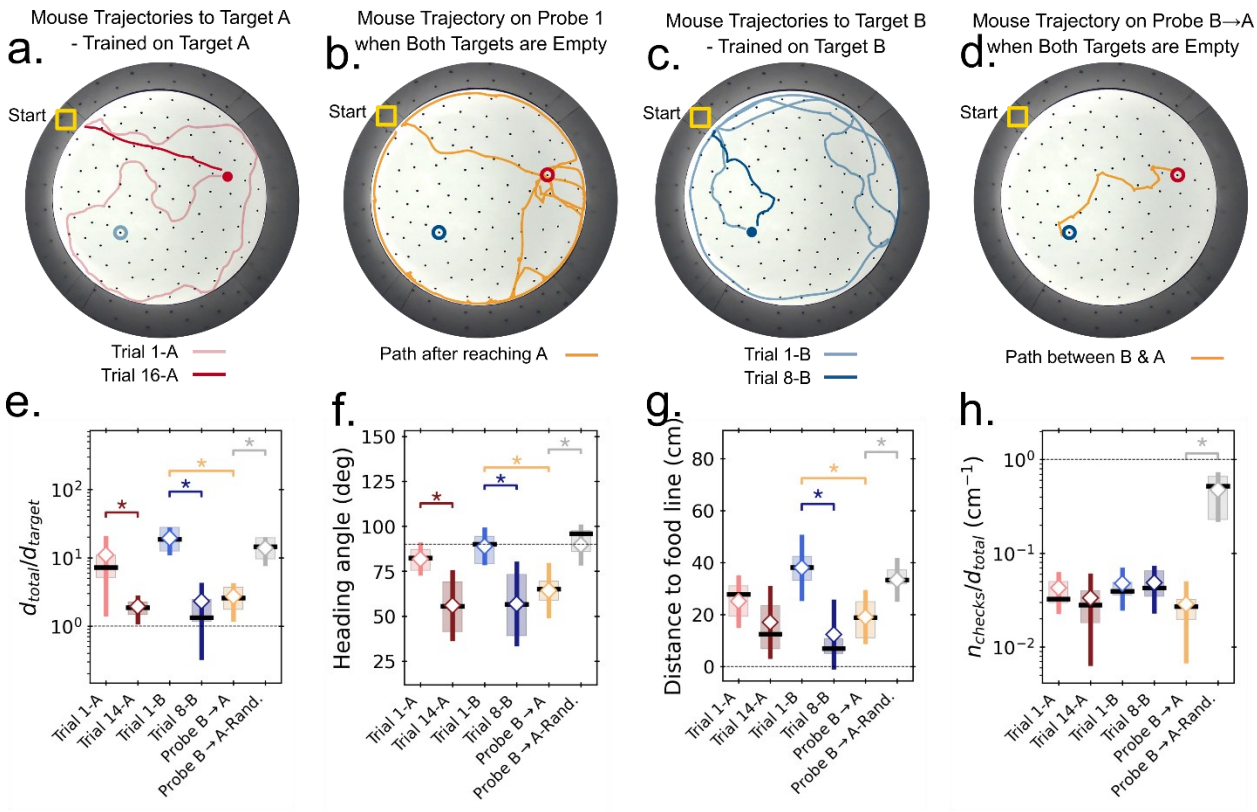

Figure 5. Two Food Location Experiment. a. Mouse trajectory towards Target A (solid red circle) during early training (light red, trial 1-A) and late training (dark red, trial 16-A). Note that, on this trial, the mouse comes within 16 cm from the where food will be on Target B in later trials; the future location of Target B is indicated by an open blue circle. **b.** Mouse trajectory from Target A (trained) during first probe without training on Target B (open blue circle – no food). After failing to find food at A, the mouse initiates a random search. **c.** Mouse trajectory towards Target B during early training (light blue, trial 1-B) and late training (dark blue, trial 8-B). After training the mouse again achieves a near direct trajectory to Target B. **d.** Mouse trajectory between Target B (trained) & Target A (trained) during second probe; no food at either location. The initial trajectory from Start to Target B is suppressed for clarity. **e.** The normalized total traveled distance d_{total}/d_{target} significantly decreases after training on Target A (Trial 1-A vs Trial 14-A) and training on Target B (Trial 1-B vs Trial 8-B). The normalized distance between Targets A and B in the probe trials are very similar to the distance between entrance and Target A (Probe B→A vs trial 14-A) and between entrance and Target B (Probe B→A vs trial 8-B). The Probe B→A distance is significantly less than the distance between two random points in the Probe B→A trial (Probe B→A trial vs Probe B→A-Rand). See below for how random points were chosen. **f and g.** Similar observations can be made for the heading angles and the distance to food line (See Extended Data for the geometric definition of these quantities). **h.** The density of checked holes

over distance traveled remains virtually unchanged for the Probe B→A compared to the learning trials 1-A, 18-A, 1-B and 8-B, and is only different from the density of checked holes in the randomly chosen sub-trajectories within Probe B→A. Stars mark significant differences ($p<0.05$). The Probe A-B - Random trial is calculated by selecting ten random pairs of points in the trial trajectory that are separated by the B to A distance; and then averaging the particular variable over these distances – see Methods for more details. In the Probe B→A trials, the food line is the straight line that connects B to A, along which the reference distance d_{target} is measured between A and B. In the other trials, the food line is a straight line from start to the specific target, either A or B, along which d_{target} is measured between start and target.

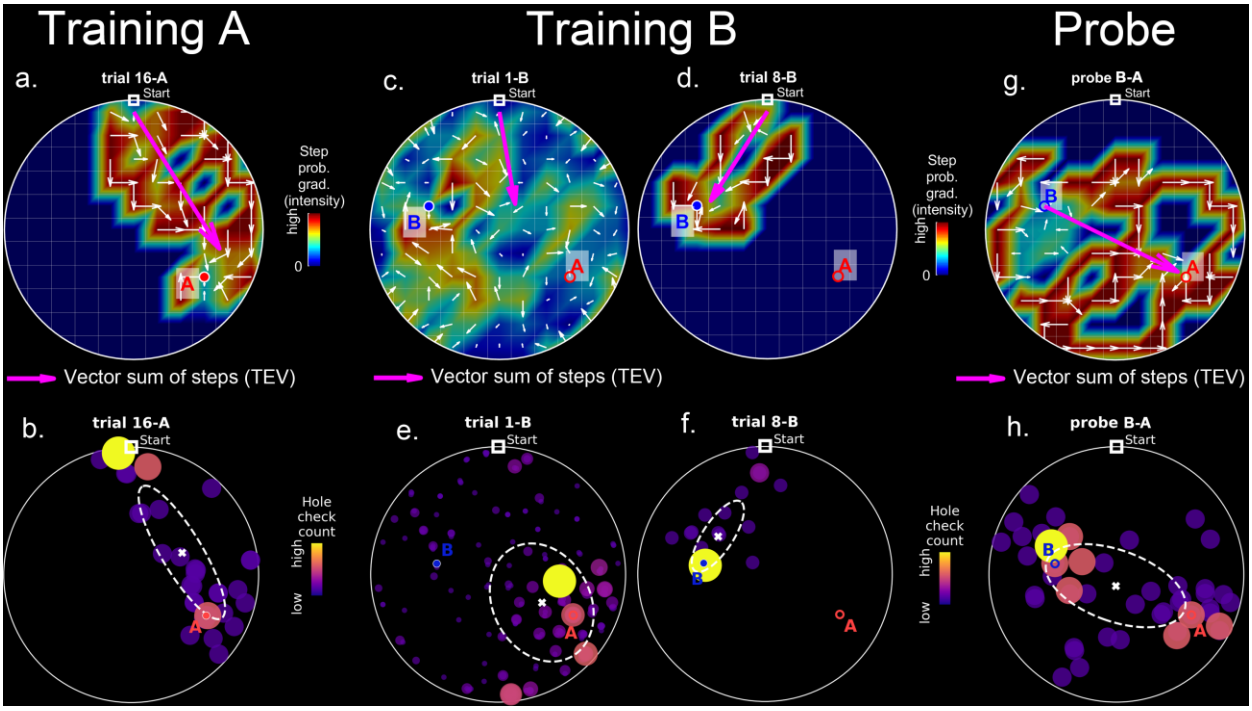

Figure 6. Step maps (top) and Active sensing (bottom) for 2-targets experiment. Each panel on the top row corresponds to the one immediately below it. **a.** Step map of the late learning trial for target A, showing probability flows symmetrically leaving and returning to the TEV directed to Site A. **b.** Active sensing for a late learning trial in target A, achieving the same result as in the static entrance case. **c and d.** Step map evolution for the learning of target B having previously learned target A. Notice the initial random flows and the appearance of a new flow towards B in trial 8-B, together with the TEV pointing to food after training. We show the typical mice (six out of eight that went straight to B in trial 8-B). Only two out of eight mice visited target A before going to B; we ignored these two mice, since the visits to A were sporadic. **e and f.** Active sensing profile evolution over the learning trials for target B for mice that had been trained first in target A (panel b). Hole-checking events accumulate near target A in the first training trial of target B (panel e), but then spread over to B (panel f). **g.** Step map for the Probe B-A trial, indicating a marked flow from B to A. The step probabilities again depart and return to the TEV that directly points from B to A, highlighting the short cut paths. The mice were able to take the B→A route without needing an external reference. We show the five out of eight mice that performed the route Start→B→A. The remaining did Start→A→B (two mice), and Start→B→Start→A (one mouse) – see Extended data Figure 5. **h.** Active sensing profile in the Probe B→A trial considering the part of the trajectories only between the first visits to the A and B regions (the orange trajectories in Extended data Figure 5). The hole-checking events accumulate on and in between A and B targets, showing that mice remember both locations and can infer the location of the other target after visiting one of them. **Active sensing**: the size and color of each circle indicates

the normalized frequency that the hole in that location was checked (larger and lighter yellow indicates higher frequency). **Step maps**: the color and arrows indicate the most probable route taken from a particular location (red means more probable; blue indicates equal probability of heading to any direction) – see Step Map Calculation in Methods, and Extended data Figure 7 for more details.

Materials and Methods

Animals

All animals were housed in the University of Ottawa Animal Care and Veterinary Services (ACVS) facility. C57Bl/6 wild-type male and female mice were ordered from Charles River, arriving at 8-9 weeks old. Mice were individually housed in 12h light/12h dark cycles (lights on at 11:00PM EST, lights off at 11:00AM EST). Animals had had food and water available *ad libitum*. The temperature of the room was kept at 22.5°C and the humidity was 40%. Mice were habituated in ACVS facilities for 1 week and began testing when they were 10-11 weeks old. Testing of each cohort took place over the course of approximately 2 weeks, 1 hour after the lights turned off. Subjects weighed 22-27 grams at the start of behavioural training. Both male and female mice were used; the same results were obtained in both sexes and were therefore pooled. All animal procedures were conducted with the approval of the University of Ottawa's Animal Care Committee and in accordance with guidelines set out by the Canadian Council of Animal Care.

Apparatus Design & Setup

The Hidden Food Maze (HFM) is a framework that trains mice to search for a food reward hidden in an open, circular arena (**Figure 1**). The protocol for the task was inspired by the Cheese Board spatial task in which rats are placed inside an arena with many holes in the ground, one of which contains a food reward[26]. Unlike the Kesner et al task, the HFM does not require the animal to be handled, and the pattern of holes is arbitrary rather than grid-like. This task design was also chosen to mirror the setup of the Morris Water Maze (MWM), which can test allocentric or egocentric navigation. Like the MWM, mice are searching for a target location in a circular environment after starting from one of four starting home locations. External landmarks can, if desired, be placed on the maze walls or local cues placed in the maze. In contrast to the MWM, this spatial learning task is a dry maze that is food motivated, which is more naturalistic and less stressful than motivation by the aversion to swimming. Notably, the task has specific design features to control for against unwanted cues, such as odours, visual cues, and handling.

The maze has a removable floor that can be washed & rotated between trials to eliminate odor trails which mice from previous trials might leave behind. The circular floor is 120 cm in diameter and has 100 holes (1.2cm diameter) arbitrarily dispersed throughout the surface, with 25 holes in each quadrant. The distance between each hole is, on average, 10 cm. The pattern of the holes is rotationally symmetrical, so the pattern looks the same regardless of whether the floor is rotated 90°, 180°, or 270° with respect to the mouse's entrance. This ensures that the mouse will have the same initial view of the holes regardless of which entrance it starts from and how the floor is rotated. Each hole is encircled by a 1cm plastic rib, sticking downwards, which can be

capped at the bottom to hold food that remains invisible from the surface. Thus, the mice cannot discern the contents of the hole just by looking across the floor from their home, but instead need to approach the hole and look inside. The surface of the floor is sanded to be matte to avoid generating reflections from the lights which might interfere with the cameras or distract the mice.

The maze floor is circular and uniform from the inside so as not to provide any directional cues. The floor of the arena is encircled by tall black walls. The walls are made of solid black PVC plastic which forms a cylinder around the maze and is open at the top. The walls are 1 cm thick and 45 cm tall, which is tall enough so the mice cannot jump out. The walls are symmetrical and designed to eliminate geometric cues that would give away directional information from asymmetries in the environment shape. Mice can use odors as landmarks for spatial learning[9] and we wanted to eliminate local food odour from a filled hole as a landmark. The arena is resting on a subfloor that contains crumbled food; food odor will diffuse through the open holes and saturate the maze thus masking the odor from a food-filled hole.

The home cages attach to the main arena by being slotted into each entrance. The home cages are 27cm x 16.5cm x 45cm and are open at the top. They contain a food hopper and a water bottle feeder. The dimensions of the home cage are based on commercial mouse cages and comply with the Canadian Council on Animal Care mouse housing standards.

When moving mice to a different starting quadrant, the home cages are designed to slide out from the maze and into a new entrance so the mice do not need to be handled and therefore stressed. The home cages include their own doors, separate from the doors that provide entry into the maze, so the home cage can be freely moved and the mice remain securely confined. The detachable home cages allow the mice to be moved to different starting locations, allowing us to control against navigational strategies that rely solely on response learning or path-integration from a fixed starting point.

Mice can perform the entirety of the trial without experimenter handling. The maze doors are designed to slide upwards and provide access to the main arena. When the trial is finished, the experimenter can slide the doors back into place and the mouse is confined to its home cage once again. To not disturb the mice, the doors are left open while they navigate.

Four LED flashlights were aimed at a white ceiling in order to create dim, diffuse lighting throughout the maze. The illuminance is measured to be 50 lux at the surface of the arena. Dim light was used to allow the mice to properly see all visual cues as well as the maze details. This

procedure is assumed to not perturb the nocturnal cycle of mice[27, 28]. Black curtains surrounded the maze to prevent the interference of non-controlled visual cues.

We used additional experiments to control for possible visual or odorant cues from the open home door. We did some experiments in total darkness using infrared LEDs with emission spectra detected by our camera. In order to exclude any potential olfactory cues emanating from the open cage door we did experiments where odor absorbent kitty litter was placed on the cage floor and two connected exhaust fans (AC Infinity MULTIFAN S5, Quiet Dual 80mm USB Fan) placed above the home cage. At their lowest (quiet) setting, the fans drew air from the home cage and blew it to the outside; replacement air from the outside came in via small openings at the bottom of the home cage. The home cage air volume was evacuated 30 times/minute, effectively equilibrating the cage and open maze air before the doors were slid upwards. Fans were turned on for 5 minutes (150 evacuations) before the trial started and for the duration of the trial. After the door was opened, the fans induced negative pressure to eliminate diffusion of odors from the home cage to the maze. The two fans were placed over each home cage to eliminate a potential directional noise cue.

Behavioural Training

Pre-training

Five days before training, mice are transferred from standard animal facility cages to the experimental home cages for habituation. On the first day, mice are allowed to habituate to their new home cages. On the second day, the door dividing the home cage and the arena is removed and each mouse is given free access to explore the empty arena for 10 minutes. There are no extra-maze landmarks present during pretraining. Animals are food restricted, with 10% of their body weight in food given back each day which they could consume *ad libitum*. On day 3, randomly chosen 50% of the holes in the arena are filled with food treats (a piece of Cheerio), and each mouse is given 10 minutes to forage for food. This is repeated on day 4, where 25% of the arena's holes are filled with food treats. On day 5, only four holes placed at the maze center contained treats. Mice pass the pre-training stage when they successfully find treats at all four locations within 20 minutes. Mice mostly confined their search to circular trajectories near the wall of the maze for the first two days but, after training with food near the maze center, they checked holes throughout the maze (**Methods Figure 1**). Training in the spatial learning task then commenced.

Visual Cues + Randomized Entrances Training

The mice are trained to locate a food reward that has a fixed relationship with four visual cues on the walls. The protocol mimics classic Morris Water Maze setups to test allocentric landmark-

based learning. One of the holes in the area is capped from the bottom with a food-containing insert that is not visible from the surface. Extra-maze landmarks (described in the Visual Cues section) are placed on the walls of the maze to serve as location cues. At the start of the trial, the door is slid upwards so the mice can enter the maze, and the mice are given a maximum of 20 minutes to find the target hole. The home door is kept open to permit the mice free access to return to their home cage during the trial; preliminary experiments indicated that closing the door perturbed the mice. The door is re-inserted after the mouse has found the food reward and returned home. If the mouse has not returned home by itself after 1 minute of finishing the food reward, it is gently guided back by the experimenter. There is an inter-trial interval of approximately 20 minutes.

Initial experiments used two trials per day but this was increased to speed up learning; our analyses and graphs are truncated at 14 trials to permit averaging over all the mice. In most experiments, three trials a day were given, over six days; at the end of each trial the mouse's home is moved to a different entrance. Preliminary experiments showed that the mice could learn the food location if the home was only moved at the end of each day. The location of the entrance is randomized with the following constraints: no two trials in a row have the same entrance, and all entrances are selected the same number of times so there is no bias. On the 7th day, a probe trial is given where the mouse is allowed to search for food in the absence of any food. Learning continues over 1 more day and, on the 9th day, a reversal trial is given where the visual landmarks are rotated by 180°. The location of the home cage is randomized before every trial. In this task, the mouse would have to learn the invariant relation between the food hole and up to four visual cues in order to locate the food.

Visual Cues: Four black symbols on white backgrounds: a square, a cross, vertical bars, and horizontal bars. The cues are taped 5 cm above the floor. The square was 15 by 15 cm. The cross consisted of two 2.5 cm wide and 14 cm long bars. The four vertical and four horizontal bars were 2.5 cm wide and 14 cm long. Previous behavioral studies have shown that mice can discriminate the visual stimuli we use[29-31]. Recording of neurons in mouse V1 have additionally shown that orientation selective cells in mice visual cortex can discriminate our visual stimuli[32].

Visual Cues + Static Entrance Training

This protocol allows mice to navigate by potentially using visual cues in cooperation with path integration or by path integration alone. The setup is the same as visual cues training, except mice enter from the same entrance each trial instead of a randomized entrance. Mice are considered well trained once their latency learning curve has plateaued for 3 successive trials. The mice's latencies had plateaued after 14 trials but training continued till 18 trials.

After 18 trials, reversal trials were done with no food present and the mouse's home cage rotated 180°. The mouse was allowed to search for food for 10 minutes. If the mice were able to learn the invariant landmark/food spatial relationship, we would expect them to search at the learned food site. If they had used a path-integration learning strategy to navigate from the fixed entrance to the food, we would expect them to search at the rotationally equivalent location (REL).

Control Experiments

Rotations following Static Entrances – with and without visual cues.

We performed additional tests for the use of visual landmarks in the Visual Cues + Static Entrance protocol. Well-trained mice had their home cages rotated to another quadrant and tested on how well they found the food from a new starting location. Mice were rotated 180°, 90°, and then 270° with respect to their original location (Fig. 3j,k,l). Mice were given 6 consecutive trials at each new location. One control group of mice had the visual cues on the arena wall removed during rotation training. If mice were able to use visual cues, we would expect them to improve their search efficiency when visual cues were available. Lack of improvement or similar performance compared to the control group without visual cues suggests that the mice do not rely on the type of visual cues we provided for spatial learning and navigation.

Navigating and Training without visual cues and in Darkness

To control for the effects of all visual information, one cohort (4 mice) was trained and tested without any cues (**Extended data Fig. 3a,b**). A second cohort (4 mice) was trained in the light according to the Visual Cues + Static Entrance protocol, then the lights are turned off and the mice are given a trial in darkness (**Extended data Fig. 3c,d**). The lights were opened in between trials so the mice did not acclimatize to the darkness. Mice were tracked using infrared LEDs with emission spectra detected by our camera. The following conditions were tested: darkness during a regular trial with food present in the arena and during a probe trial where there is no food present in the arena; both conditions gave the same results.

We additionally both trained and tested a cohort (4 mice) in darkness and with control of potential odors. A recent study has shown that placing sighted mice in darkness impairs entorhinal cortex head direction cell tuning[8] , but here the mice successfully learned to find food and with the same time course of learning (**Extended data Fig. 3e,f**).

Two Food Location Training

This protocol aims to test flexibility of spatial learning using only path integration (N=8 mice). No visual cues are placed on the arena walls. Mice entered the arena from the same entrance each

trial. Food was placed in a target well in "Target A" and mice were trained to find this location for 18 trials, 3 trials a day. A probe trial was done after trial 18. For trials 19-26, the food was moved to a different location, "Target B", and mice are trained to find the new location. We were careful to choose Target B so that Target A and B were not symmetric with respect to the mouse's entrance. A second probe trial was done after trial 26. For some cohorts, the training continued until trial 34 and a third probe was given. We additionally carried out the same analysis (N=4 mice) but with a 180° rotation of the home cage before the probe trial.

Behavioural Analysis

Path tracking was done with Ethovision XT15 (Noldus) based on contrast detection. Mice were tracked according to 3 body points at 30 frames per second on a 1080P USB camera with OV2710 CMOS. Videos were first recorded on AMCap webcam recording software and imported into Ethovision in order to preserve high quality videos. Trajectory plotting was done in Python.

Latency to target was calculated from when the nose point of the mouse enters the arena until the nose of the mouse enters the target hole. Trajectory analyses are based on the nose point sequence of the mouse. Speed and distance were calculated based on trajectory coordinates exported from Ethovision.

Search bias during probe trials is calculated by totalling the time a mouse spent within a 30x30cm square centered around the target vs the REL target in the other 3 quadrants during a 2-minute trial.

Statistical Tests

Significance testing was conducted in a manner that was appropriate for each dataset. Paired comparisons utilized the t-test for parametric data and the Wilcoxon signed-rank test for non-parametric data. The significance cut-off was set at 0.05. Statistics were calculated using R: A language and environment for statistical computing. URL http://www.R-project.org/.

Analysis of Trajectories

Experiment Alignment

The procedure described in this section is applied to calculate the Step Maps (and hence the target estimation vector, TEV) and the Hole Checking Spatial Distributions. Each mouse in the Random Target experiment entered from a different quadrant of the arena, for each consecutive trial, and the target was fixed relative to the arena (global reference frame). However, in order to increase sampling, we need to coherently align the mice entrances, generating the mouse's perspective reference frame (see **Methods Figure 2**). Each experiment batch was performed with

four mice, such that at any given trial, a given mouse entered from quadrant 1 (**Figure Methods 2b**), another entered from quadrant 2 (**Methods Figure 2c**), the next entered from quadrant 3 (**Methods Figure 2d**), and the last entered the arena from quadrant 4 (**Methods Figure 2e**). Most of the experiments, then, consisted of two batches of four mice that had to be conveniently aligned to increase sampling. In this figure, we aligned all the entrances to the Start point at the top of the arena (trajectories were rotated accordingly), emphasizing the four possible positions of the target from the mouse's perspective.

We had to find a way to align all the targets in a given trial to one of the four possible positions, such that the experiment stays coherent over trials (i.e., the target randomly switches in between these four positions from trial to trial). For example, in **Methods Figure 2 panels b-e**, we show trial 14. The target position for each starting quadrant is labeled with a red letter A and a red circle, whereas the target for the previous trial is labeled with a green circle and a green letter A subscripted with "trial 13". These two positions (trial and previous trial) must always be different to keep the random characteristic of the experiment.

Notice that the target positions in each of the c, d and e panels in **Methods Figure 2** is simply rotated relative to the target positions in panel b. The target configuration in panel c is 90° clockwise rotated relative to the target configuration in panel b. The configuration in panel d is 180° clockwise rotated, and panel e, 270° clockwise rotated; both angles are relative to the configuration in panel b. Thus, in order to sample all the mice together, we simply rotate the trajectories from the mice that started in quadrant 2 (panel c) counter-clockwise by 90°; the trajectories from the mice that started in quadrant 3 (panel d) are rotated by 180° counter-clockwise; and the trajectories starting from quadrant 4 (panel e) are rotated by 270° counter-clockwise. This reduces all the experiments to the first quadrant, allowing us to sample all the trajectories of all the mice together in each individual trial.

Active sensing and hole-check detection

The active sensing detection was performed by an automated algorithm that captured the attempts of the mouse to check a hole for food. Given the arena design, the mouse needs to put its nose very close to the hole to be able to see the food or detect any food odour. With that in mind, we defined two methods for detecting a hole check (see **Methods Figure 3**, and Supplementary Video). The two methods are complementary and the second can find hole-check events missed by the first method: we apply the first method, then perform a visual check on the data to see if there were potential hole-checks that were missed. Then, we apply the second method to capture any remaining events.

The first method is the minimum velocity criterion and provides a high threshold for hole check identification. Four conditions must be satisfied simultaneously to detect an event: (i) the nose of the mouse must be within 3 cm of an arena hole; (ii) the velocity has to be less than 20% its maximum value; (iii) the velocity must be at a minimum; and (iv) the velocity has to have dropped by at least 5 cm/s to reach that minimum.

The second method is more inclusive and defines a hole-checking event by a simple slowing down event, provided that: the nose is within 3 cm of an arena hole, and the slowing down is enough for the velocity to drop past the 20% threshold of its maximum.

The detected events by the application of these two methods in sequence are marked for all trajectory samples of the Probe B→A in **Extended Data Figure 5** as examples. A particular case is shown in more details in **Methods Figure 3**, and in the **Supplementary Information Video** (mouse_36_Probe2_hole_check_synchronized.mp4**)** with the recording of a mouse's performance. It is worth mentioning that in the video, there is a clear hole-check miss by our algorithm in the 25 to 27 second range. This is because the velocity of the mouse was already below the 20% maximum before and after the hole check, hence the two sets of criteria defined above were not met for this particular event. Otherwise, we clearly see that all the other events are captured by sequential use our two methods.

Step map calculation and Target Estimation Vector (TEV)

The step maps are a way to visualize mouse trajectories and directionality across all the mice for each trial. They are related to a velocity map of the arena, except that in the Step Map, the arrow points in the direction of most probable movement (instead of the direction of the velocity). **Methods Figure 4** shows a scheme for how we compute this map, and we detail it below. The procedure here is applied to each individual trial independently (either during learning, or for a probe trial). Box dimensions in Figure 4 are calculated with methods described in Falconer[33].

We start by overlaying a lattice of size L on top of the arena recording. This means that there are L boxes on the x (horizontal) direction, and L boxes on the y (vertical) direction, making a total of L^2 boxes in the lattice. Each box defines a lattice site with coordinates (x, y), such that x and y are integers between 1 and L. Each mouse corresponds to an independent observation of a trajectory, so we overlay all mice trajectories for each trial in the calculations below. Next, we map the coordinates of the trajectories into the lattice coordinates in order to obtain a temporal sequence of visited lattice sites in each trial. We are interested in counting the number of times that each subsequent pair of adjacent lattice sites appears in this sequence, regardless of what mouse it came from. This will be used to define the step map, detailed in what follows.

The step map $\vec{M}(x,y)$ is a vector field defined on the lattice that assigns to each site the preferential direction of movement. This direction is directly estimated from the trajectories data. Mathematically, it can be written as $\vec{M}(x,y) = [M_h(x,y) \;\; M_v(x,y)]$, where M_h is the horizontal component and M_v is the vertical component. The horizontal component expresses the trend to move from left to right (along the lattice x-axis), and the vertical component expresses the trend to move in the perpendicular direction (*i.e.*, "up" or "down" in the top view of the lattice of Methods Fig. 3a, along the lattice y-axis). Each component is a function of the lattice coordinates, (x,y). The vector $\vec{M}$ is defined as the spatial gradient of the probabilities to move out of (x,y) towards one of its adjacent sites. Thus, its components are given by

$$\begin{aligned} M_h(x,y) &= P_{\rightarrow}(x,y) - P_{\leftarrow}(x,y)\,, \\ M_v(x,y) &= P_{\uparrow}(x,y) - P_{\downarrow}(x,y)\,, \end{aligned} \qquad \text{Eq. (1)}$$

where $P_{\rightarrow}(x,y)$ is the probability of stepping right, *i.e.*, from (x,y) to $(x+1,y)$; $P_{\leftarrow}(x,y)$ is the probability of stepping left [from (x,y) to $(x-1,y)$]; $P_{\uparrow}(x,y)$ is the probability of stepping up [from (x,y) to $(x,y+1)$]; and $P_{\downarrow}(x,y)$ is the probability of stepping down [from (x,y) to $(x,y-1)$]. The probabilities of stepping out of (x,y) must be normalized, therefore $P_{\rightarrow}(x,y) + P_{\leftarrow}(x,y) + P_{\uparrow}(x,y) + P_{\downarrow}(x,y) = 1$. If there is no trajectory going through a particular site, the probability of going from that site into each of the directions is equal to the "null" probability, $P_0 = 1/4$.

The stepping-out probabilities are computed from the mouse trajectory in the following way. This procedure is applied to all mice overlayed together in each individual trial. First, a step is defined as moving from a box at (x,y) to one of its four adjacent boxes, say the one on the right $(x+1,y)$. Now, if we wanted to calculate the probability of going right from a position (x,y), we need to go through the sequence of visited lattice sites looking for (x,y) in an instant followed by $(x+1,y)$ in the immediately next instant. We count the number of times $S_{\rightarrow}(x,y)$ that this pair appears. The count of the transitions of (x,y) to $(x+1,y)$, in this example, is made regardless of when these transitions happened during the trajectory. The only condition is that the position (x,y) must be immediately followed by $(x+1,y)$. The same counting is made for the transitions from (x,y) to $(x-1,y)$ and stored in $S_{\leftarrow}(x,y)$, from (x,y) to $(x,y+1)$ in $S_{\uparrow}(x,y)$, and from (x,y) to $(x,y-1)$ in $S_{\downarrow}(x,y)$.

With these sums of steps performed, we can calculate the probability of stepping right, left, up or down. First, note that initially the chance of going to any direction is P_0, assuming we know nothing about the trajectories. Now, given that we observed $S_{\rightarrow}(x,y)$ steps going to the right, we need to update the chance of going to the right using the union (i.e. sum) of the observed

chance $S_{\rightarrow}(x,y)/N$ with P_0, where N is the total number of steps counted toward any direction in all lattice sites in a given trial:

$$W_{\rightarrow}(x,y) = \frac{S_{\rightarrow}(x,y)}{N} + P_0 - \frac{S_{\rightarrow}(x,y)}{N} P_0 \,, \qquad \text{Eq. (2)}$$

where $W_{\rightarrow}(x,y)$ is the direction weight of the action "step to the right from (x,y)". Alternately, this can be written as a weighted sum of memory-driven stepping with probability 1 and random stepping with probability P_0:

$$W_{\rightarrow}(x,y) = \frac{S_{\rightarrow}(x,y)}{N} + P_0 \left(1 - \frac{S_{\rightarrow}(x,y)}{N}\right).$$

We employ Eq. (2) because the mouse could, in principle, have chosen any of the other three directions. This ensures that every time the memory is increased (adding $S_{\rightarrow}/N$ to the weight), the random contribution for that step decreases; this justifies the last term in Eq. (2), where the intersection between the memory term and the "null" probability, $S_{\rightarrow}P_0/N$ is subtracted from the aforementioned union. One can easily see that the formula guarantees that the null probability is automatically recovered, i.e. $W_{\rightarrow}(x,y) = P_o$, when there are zero observed steps in a given direction.

Finally, we use these weights to calculate the probabilities, first defining the total weight of stepping out of (x,y),

$$W_{out}(x,y) = W_{\rightarrow}(x,y) + W_{\leftarrow}(x,y) + W_{\uparrow}(x,y) + W_{\downarrow}(x,y) \,, \qquad \text{Eq. (3)}$$

and then calculating the probabilities by

$$P_{\rightarrow}(x,y) = \frac{W_{\rightarrow}(x,y)}{W_{out}(x,y)} \,. \qquad \text{Eq. (4)}$$

The same is made for the other directions (left, up and down). Eq. (4) ensures that the probability of stepping out of (x,y) is always normalized. This is then applied to Eq. (1) to obtain the step map.

Let us go through the simple example shown in Figure Methods 3b. In this case, the accumulated trajectories in the center box located at $(x,y) = (3,3)$ are such that there are two passes going up and one pass going down. Thus, each direction has the following weight [given by Eq. (2)]: $W_{\uparrow}(3,3) = \frac{2}{3} + P_0 - \frac{2}{3}P_0 = \frac{3}{4}$, $W_{\downarrow}(3,3) = \frac{1}{3} + P_0 - \frac{1}{3}P_0 = \frac{1}{2}$, and $W_{\rightarrow}(3,3) = W_{\leftarrow}(3,3) = P_0 = \frac{1}{4}$. Applying Eq. (4), we obtain the probabilities $P_{\uparrow}(3,3) \approx 0.43$, $P_{\downarrow}(3,3) \approx 0.29$, and $P_{\rightarrow}(3,3) =$

$P_{\leftarrow}(3,3) \approx 0.14$. This results in a step map vector $\vec{M}(3,3) = [0 \quad 0.14]$. This vector represents the preferred direction of movement out of site $(3,3)$, meaning the mice are likely to pass in the vertical direction, going from bottom to top (hence a positive vertical quantity); and are not picky regarding left or right (hence a null horizontal quantity).

The target estimation vector (TEV) is computed as $\vec{D}$, i.e., the total sum of the gradient over all sites, $\vec{D} = \sum_{(x,y)} \vec{M}(x,y)$, where the sum runs over all the (x,y) pairs in the lattice. This quantity gives the net direction and sense of the trajectories represented by the map $\vec{M}$. The magnitude of $\vec{D}$ may extrapolate the arena size in a strongly directed flow. This is because it measures the total path length represented by the map, and hence it can be of order L^2. Conversely, the arena diameter is of order L. Thus, in the trials where there is a strongly directed flow, we rescale the intensity $\vec{D}$ to match the distance to the target, $\vec{D} = d_{target}\hat{D}$, where $\hat{D} = \vec{D}/|\vec{D}|$ is the unit vector and d_{target} is either the distance from start to target (for training trials), or the distance between the two targets (in the Probe B-A trials). Otherwise, the TEV would pass over the target and extrapolate from the arena. Note that this does not change the direction and sense of this vector. In trials where there is no preferred direction, the magnitude $|\vec{D}|$ is roughly the distance from start to the arena's center, and hence it does not need rescaling.

The quantities $\vec{M}(x,y)$ and $\vec{D}$ (TEV) are plotted in **Figures 4 and 6 (top rows)**. The color of the map is the magnitude $|\vec{M}(x,y)| = \sqrt{M_h^2 + M_v^2}$. These quantities are also shown in greater detail in **Extended data Figures 4 and 6**.

Data Availability

All codes and data are available in https://github.com/neuro-physics/mouse-cogmap

Methods Figures

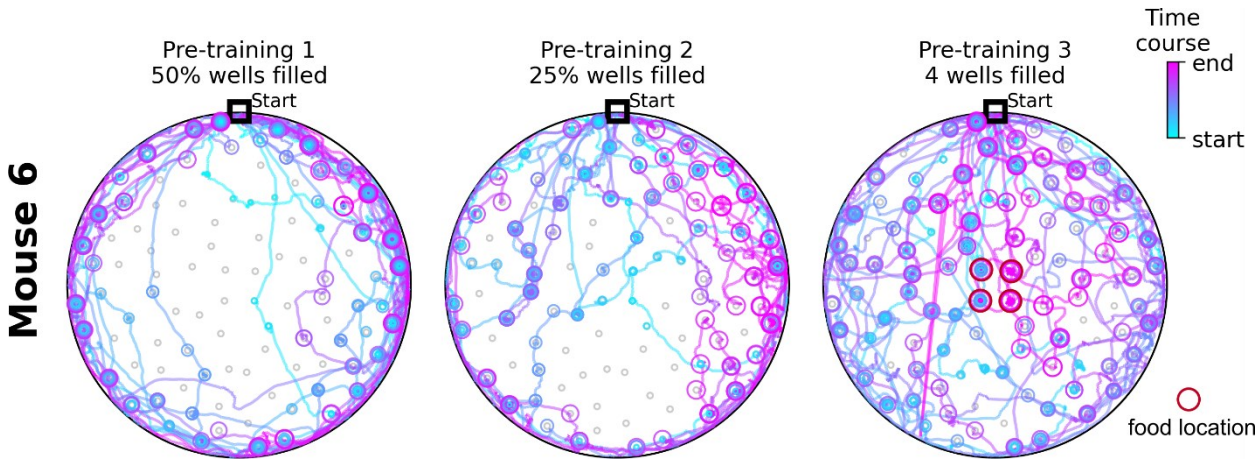

Methods Figure 1: Pre-training trajectories. Example trajectories of a mouse during pretraining. Mice starts with exploring an arena with 50% of the holes filled with food, then 25%. Most trajectories are confined near the maze walls. During the final pre-training, only the central 4 wells (red circles) have food and the mice now explored the entire maze. The mice can re-enter their home cage at any time during the food search resulting in multiple blue and magenta lines converging at the Start location.

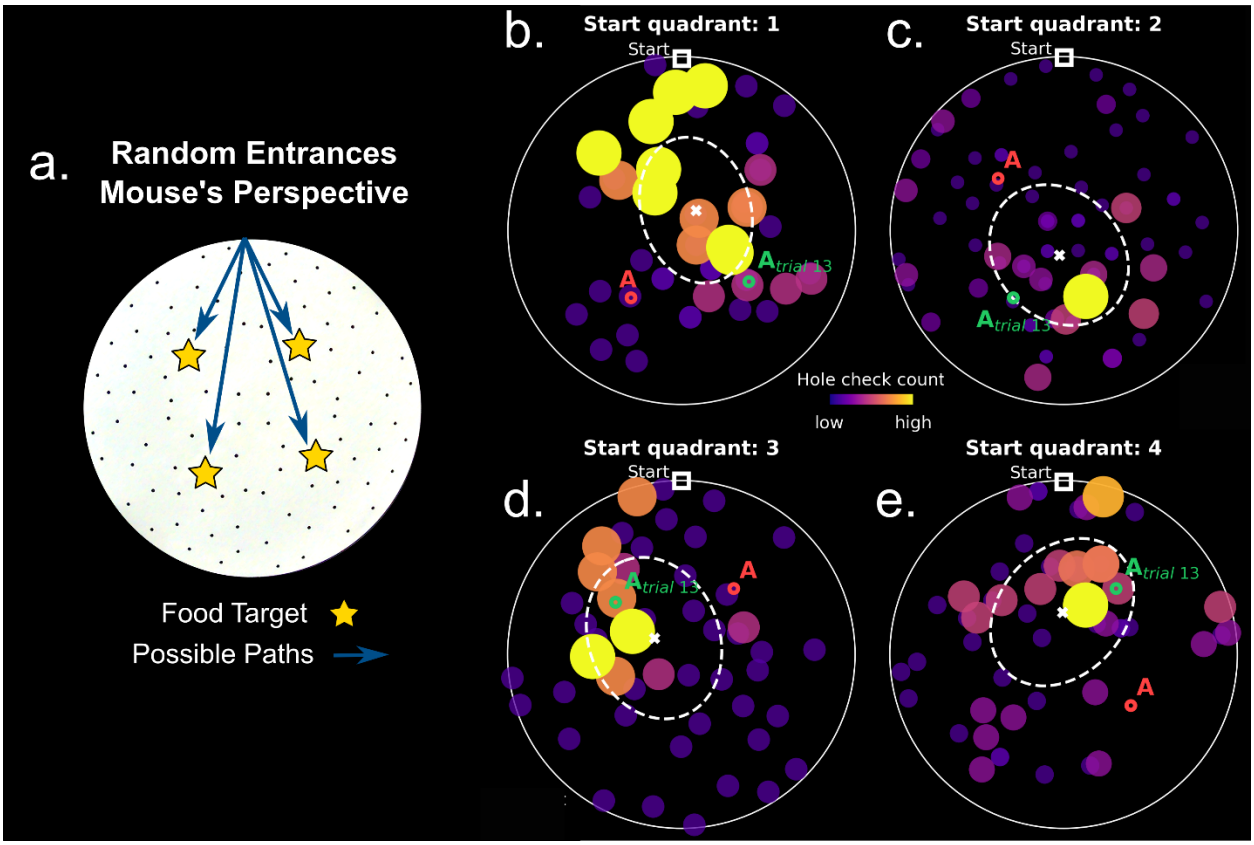

Methods Figure 2: Active sensing from the mouse perspective in random entrance experiment. a. The mouse can enter the arena from each of the four quadrants of the circle. Due to the circular symmetry, we can rotate each entrance setup so that it always aligns to the top of the screen, as shown. This is the mouse perspective of the experiment, where now the mouse has to find one of four possible target locations (labeled with a red A in panels b to e; target of the previous trial is labeled in subscripted green A). Quadrant 1 is NE (**panel b**), quadrant 2 is NW (**panel c**), quadrant 3 is SW (**panel d**), and quadrant 4 is SE (**panel e**). Moreover, notice that the configuration of targets for panel c is rotated 90° clockwise relative to panel b; panel d is 180° clockwise relative to b; and e is 270° clockwise relative to panel b. Thus, we can reverse these rotations and align all panels to the configuration in b (see Experiment Alignment in Methods). We do that consistently for every trial, increasing the sample for the active sensing of each trial. The data shown in panels b to e refer to the frequency of checks for a particular hole in the arena: number of checks divided by total number of checks, normalized between zero (purple) and one (light yellow). Larger balls represent higher frequencies.

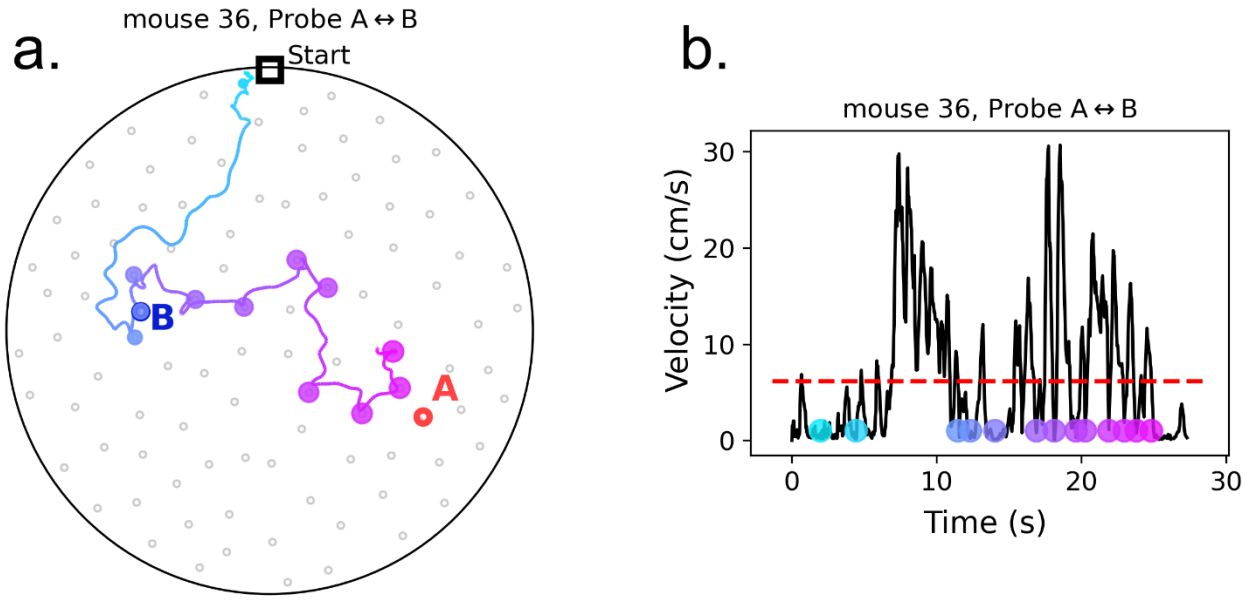

Methods Figure 3: Definition of a hole-checking event as active sensing. a. Observed trajectory with hole-checking events identified by the two criteria sets of our method, applied subsequently to avoid missed checks. **b.** The velocity profile with dots representing the hole-checking events shown in panel (a). The color from cyan to magenta represents the time order of the events, as presented in panel (a). See Hole-checking Event Detection in Methods for details of the detection criteria. A video, Methods Hole check Illustration Video mouse_36_Probe2_hole_check_synchronized.mp4, illustrates the operation of the hole check algorithm.

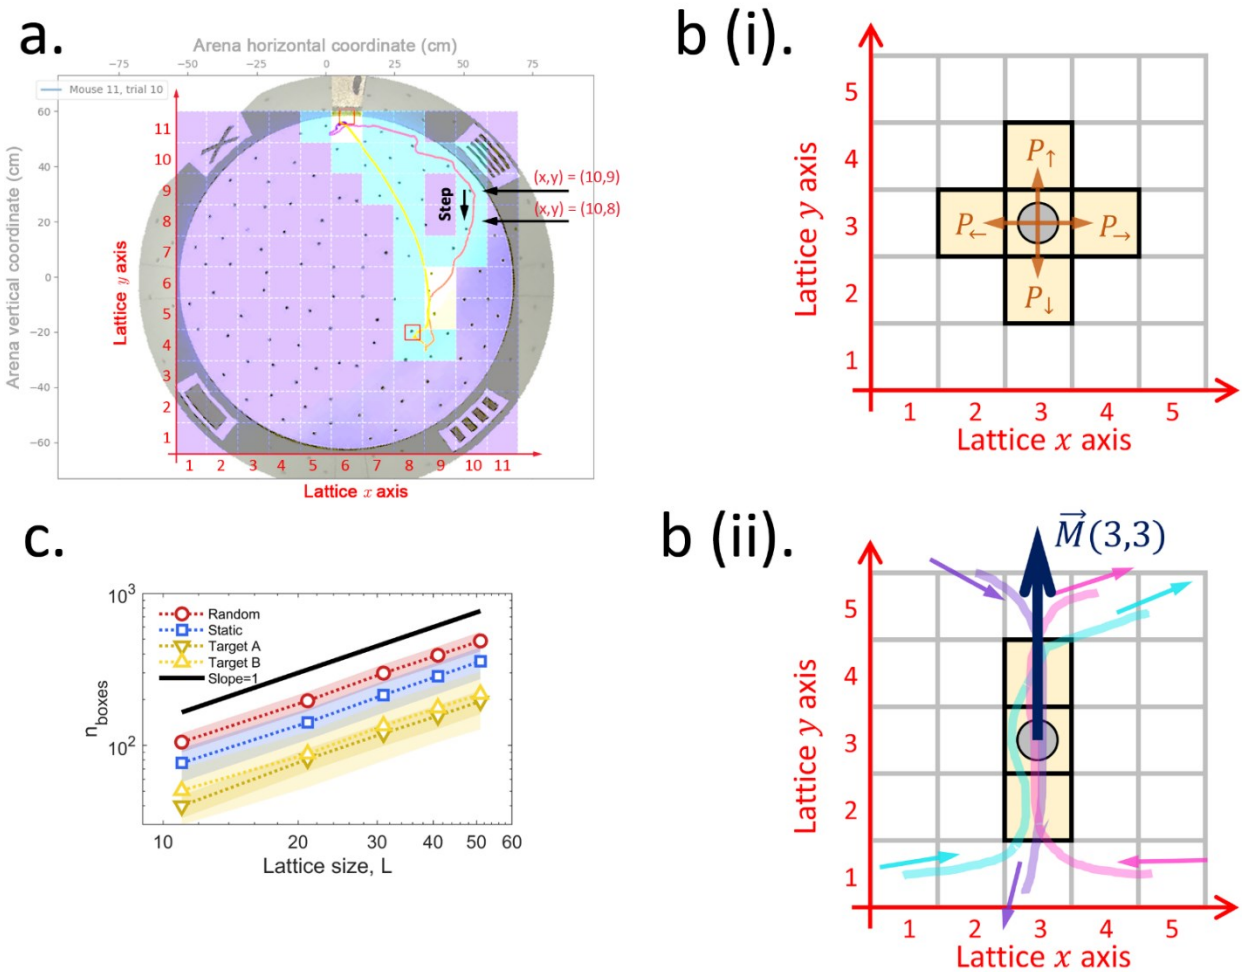

Methods Figure 4: Arena lattice, box size scaling and step map definition. a. A square lattice with L^2 boxes is overlayed on top of the arena recording to build the step map. There are L boxes along either the x or y axis. L is an odd number, such that the center box is aligned with the center of the arena. For $L = 11$, each lattice site (box) measures approximately 11 cm by 11 cm. This lattice is used to calculate the step maps (see Step Map Calculation in Methods, and panel b). **b (i).** To generate the step maps, we start by counting the number of steps a mouse takes between any two adjacent sites in the lattice; for example, when the mouse lies in the box marked by the gray circle, it can choose to take a step up, down, left or right. If no trajectories pass over that site towards a given direction, then that particular direction is assumed to have null probability $P_0 = 1/4$. **b (ii).** An example showing three trajectories that eventually pass by the marked spot, two times going up (cyan and magenta) and one time going down. The step map vector $\vec{M}(3,3)$ is then pointing up. See Methods for the detailed calculation of this vector. **c.** We count the number of boxes n_{boxes} for each trajectory (exemplified in panel a) and plot it versus the lattice lateral size L, yielding the box dimension[33] equal 1 for all trajectories in every

trial and experimental setup. The observed robust scaling allows us to use moderately small $L = 11$ for our experimental analysis; thus, larger L yields qualitatively the same results (not shown).

References

1. Chan, E., Baumann, O., Bellgrove, M.A. & Mattingley, J.B. From objects to landmarks: the function of visual location information in spatial navigation. *Front Psychol* **3**, 304 (2012).

2. Rosenberg, M., Zhang, T., Perona, P. & Meister, M. Mice in a labyrinth show rapid learning, sudden insight, and efficient exploration. *Elife* **10** (2021).

3. Collett, T.S., Cartwright, B.A. & Smith, B.A. Landmark learning and visuo-spatial memories in gerbils. *J Comp Physiol A* **158**, 835-851 (1986).

4. Eichenbaum, H., Stewart, C. & Morris, R.G. Hippocampal representation in place learning. *J Neurosci* **10**, 3531-3542 (1990).

5. Benhamou, S. No evidence for cognitive mapping in rats. *Animal Behaviour* **52**, 201-212 (1996).

6. McNaughton, B.L.*, et al.* Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *J Exp Biol* **199**, 173-185 (1996).

7. McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I. & Moser, M.B. Path integration and the neural basis of the 'cognitive map'. *Nat Rev Neurosci* **7**, 663-678 (2006).

8. Asumbisa, K., Peyrache, A. & Trenholm, S. Flexible cue anchoring strategies enable stable head direction coding in both sighted and blind animals. *Nat Commun* **13**, 5483 (2022).

9. Fischler-Ruiz, W.*, et al.* Olfactory landmarks and path integration converge to form a cognitive spatial map. *Neuron* **109**, 4036-4049 e4035 (2021).

10. Poulter, S., Hartley, T. & Lever, C. The Neurobiology of Mammalian Navigation. *Curr Biol* **28**, R1023-R1042 (2018).

11. Monaco, J.D., Rao, G., Roth, E.D. & Knierim, J.J. Attentive scanning behavior drives one-trial potentiation of hippocampal place fields. *Nat Neurosci* **17**, 725-731 (2014).

12. Engelmann, J., Wallach, A. & Maler, L. Linking active sensing and spatial learning in weakly electric fish. *Curr Opin Neurobiol* **71**, 1-10 (2021).

13. Mirmiran, C., Fraser, M. & Maler, L. Finding food in the dark: how trajectories of a gymnotiform fish change with spatial learning. *J Exp Biol* **225** (2022).

14. Grieves, R. & Dudchenko, P. Cognitive maps and spatial inference in animals: Rats fail to take a novel shortcut, but can take a previously experienced one. *Learning and Motivation* **84**, 81-92 (2012).

15. Bennett, A.T. Do animals have cognitive maps? *J Exp Biol* **199**, 219-224 (1996).

16. Savelli, F. & Knierim, J.J. Origin and role of path integration in the cognitive representations of the hippocampus: computational insights into open questions. *J Exp Biol* **222** (2019).

17. Chapillon, P. & Roullet, P. Use of proximal and distal cues in place navigation by mice changes during ontogeny. *Dev Psychobiol* **29**, 529-545 (1996).

18. Hebert, M., Bulla, J., Vivien, D. & Agin, V. Are Distal and Proximal Visual Cues Equally Important during Spatial Learning in Mice? A Pilot Study of Overshadowing in the Spatial Domain. *Front Behav Neurosci* **11**, 109 (2017).

19. Rogers, J., Churilov, L., Hannan, A.J. & Renoir, T. Search strategy selection in the Morris water maze indicates allocentric map formation during learning that underpins spatial memory formation. *Neurobiol Learn Mem* **139**, 37-49 (2017).

20. Youngstrom, I.A. & Strowbridge, B.W. Visual landmarks facilitate rodent spatial navigation in virtual reality environments. *Learn Mem* **19**, 84-90 (2012).

21. Shamash, P.*, et al.* Mice learn multi-step routes by memorizing subgoal locations. *Nat Neurosci* **24**, 1270-1279 (2021).

22. Zhu, S., Lakshminarasimhan, K.J., Arfaei, N. & Angelaki, D.E. Eye movements reveal spatiotemporal dynamics of visually-informed planning in navigation. *Elife* **11** (2022).

23. Landau, B., Gleitman, H. & Spelke, E. Spatial knowledge and geometric representation in a child blind from birth. *Science* **213**, 1275-1278 (1981).

24. Landau, B., Spelke, E. & Gleitman, H. Spatial knowledge in a young blind child. *Cognition* **16**, 225-260 (1984).

25. Banino, A.*, et al.* Vector-based navigation using grid-like representations in artificial agents. *Nature* **557**, 429-433 (2018).

26. Kesner, R.P., Farnsworth, G. & Kametani, H. Role of parietal cortex and hippocampus in representing spatial information. *Cereb Cortex* **1**, 367-373 (1991).

27. Kronfeld-Schor, N.*, et al.* Chronobiology by moonlight. *Proc Biol Sci* **280**, 20123088 (2013).

28. Upham, N., S. & Hafner, J.C. Do nocturnal rodents in the great basin desert avoid moonlight? *Journal of Mammology* **94**, 59-72 (2013).

29. Jacobs, A.L.*, et al.* Ruling out and ruling in neural codes. *Proc Natl Acad Sci U S A* **106**, 5936-5941 (2009).

30. Prusky, G.T. & Douglas, R.M. Characterization of mouse cortical spatial vision. *Vision Res* **44**, 3411-3418 (2004).

31. Prusky, G.T., West, P.W. & Douglas, R.M. Behavioral assessment of visual acuity in mice and rats. *Vision Res* **40**, 2201-2209 (2000).

32. Long, M., Jiang, W., Liu, D. & Yao, H. Contrast-dependent orientation discrimination in the mouse. *Sci Rep* **5**, 15830 (2015).

33. Falconer, K. *Fractal Geometry: Mathematical Foundations and Application* (John Wiley and Sons, 2004).

Extended Data

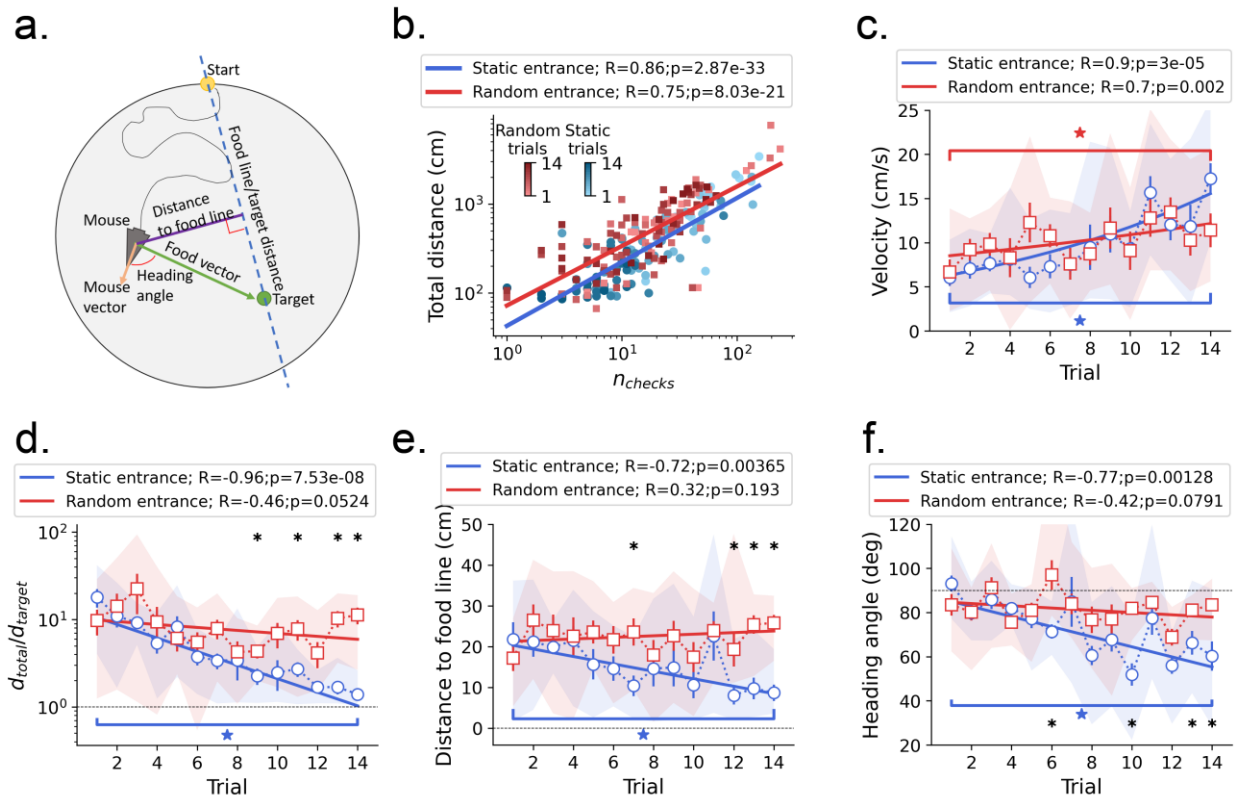

Extended data Figure 1: Geometric features and correlations with active sensing. a. Definition of the geometric features that measure performance. Food line: the straight line that connects the food hole (target) to the entrance (this line defines the optimal trajectory and minimum distance from start to target, d_{target}). Distance to food line: minimum (perpendicular) distance from the mouse center to the food line. Food vector: a vector that points from the mouse center to the target. Mouse vector: a vector that points from the mouse center to its nose. Heading angle: angle between the food and mouse vector. **b.** The total traveled distance d_{total} correlates with the number of checked holes n_{checks} for every trial. **c.** Velocity increases significantly in both static and random experiments. **d, e, f.** Comparison of the evolution of kinetic and geometric parameters across trials. The traveled distance d_{total} (panel d) becomes quasi-optimal in the static entrance case, approaching d_{target}; The deviation from the optimal trajectory (panels e and f) decreases significantly only in the static case. Asterisks mark significant differences between random and static cases (q<0.05, FDR-corrected); stars mark significant differences between first and last trials of the same case (p<0.05).

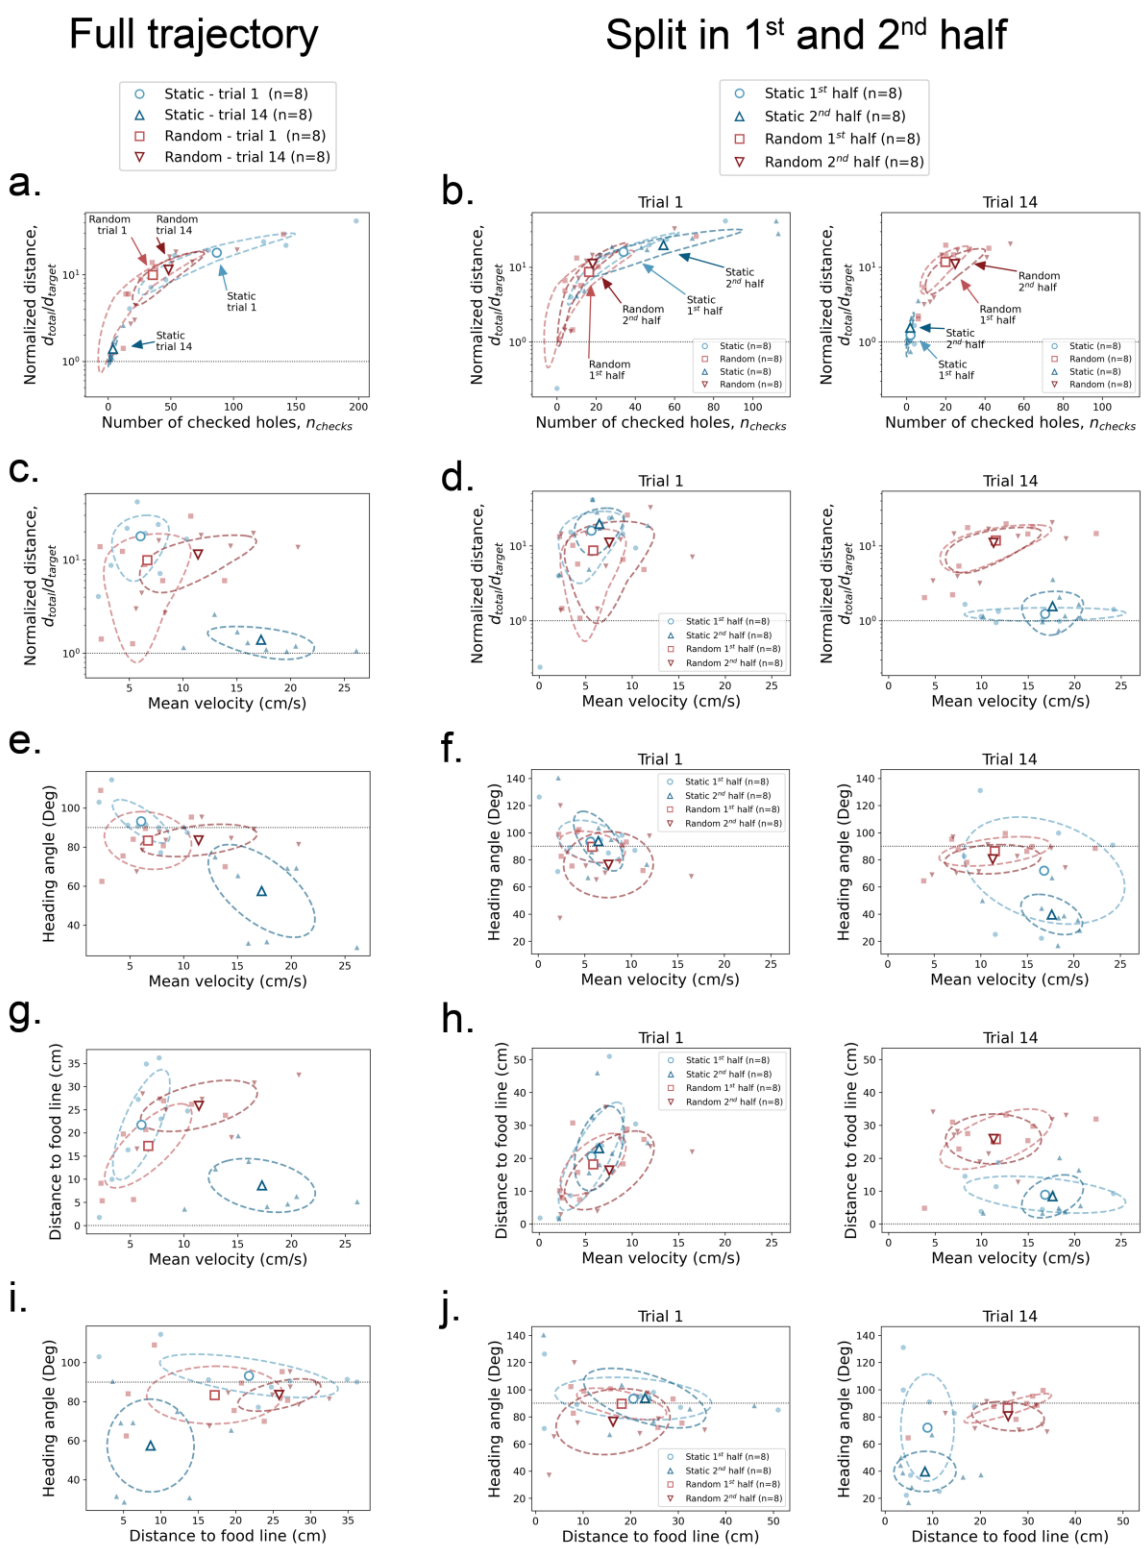

Extended data Figure 2: Covariance between geometric, kinetic features, and active sensing. Left column: full trajectory analysis – each feature is averaged over the whole trajectory of each mouse. **Center and right columns**: half trajectory analysis – each feature is averaged over the first

and second half of the trajectory for trial 1 (center) and trial 14 (right). The definition of each of these parameters is given in Extended Data Figure 1. The four large empty symbols give the average over the eight mice for each case. The dashed ellipses around them show the covariance of the features for each case. Horizontal dotted lines are there for reference values.

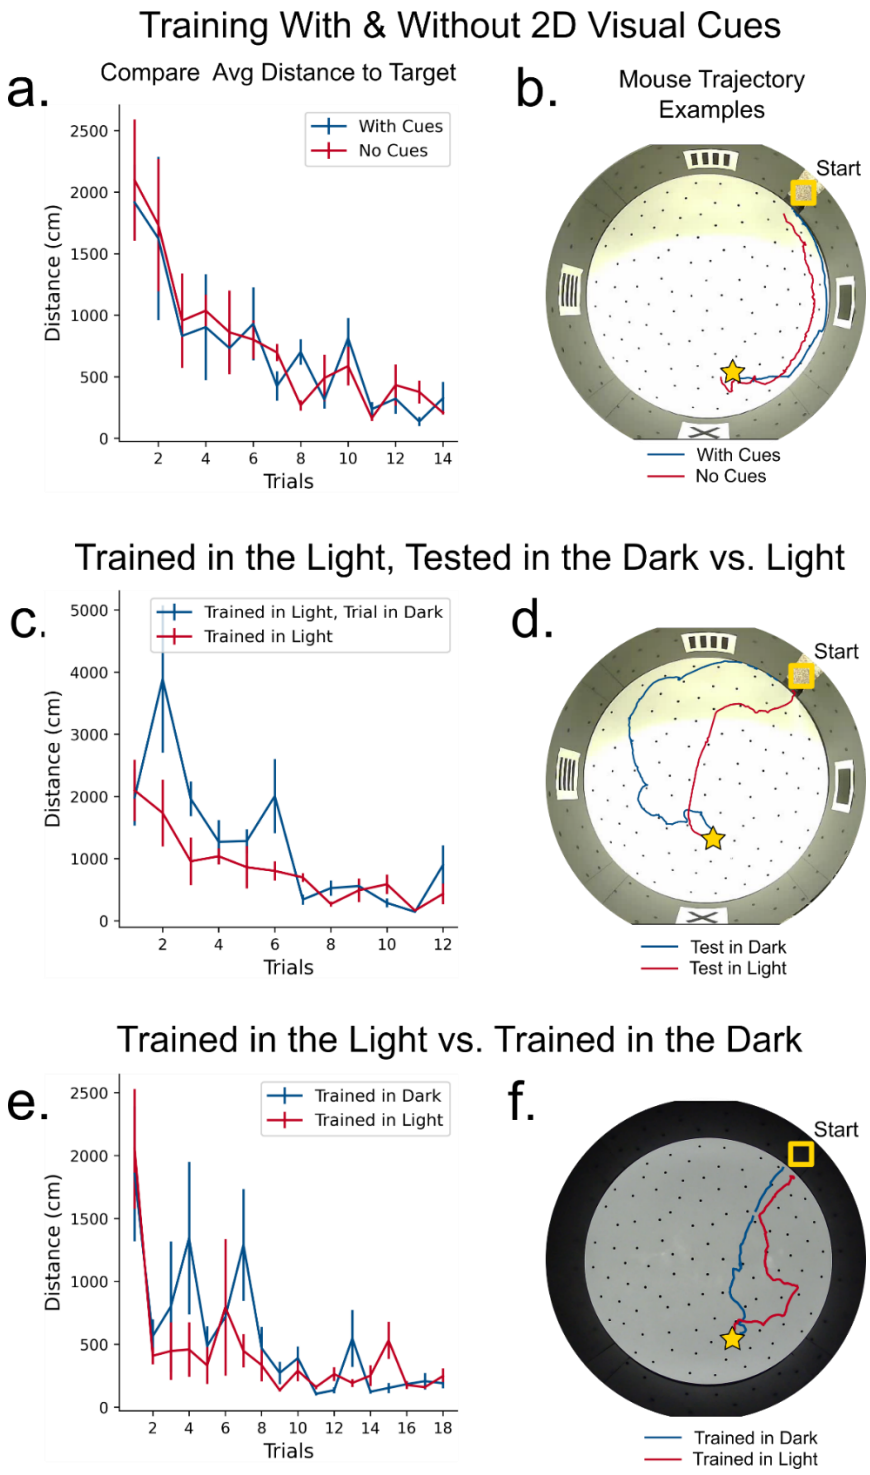

Extended data Figure 3: Control experiments for path-integration. a. Comparison in learning rate between mice with 2D cues on the wall (blue) and mice with no cues (orange). N = 4 mice per group, Error = SE. No significant difference in performance. **b.** Example trajectories of mice

trained with and without 2D visual cues. **c.** Comparison in learning rate between two groups of mice trained in the light. N = 4 mice per group, Error = SE. No significant difference in performance. **d.** An example trajectory of a mouse trained in the light successfully navigating towards the target in darkness (blue line). A control trajectory of a mouse navigating in the light is shown (red line). **e.** Comparison in learning rate between mice trained and tested in complete darkness versus mice trained in light. The trained and tested in darkness mice had additional controls for potential odor cues (see Methods). N = 4 mice per group, Error = SE. No significant difference in performance. **f.** Example trajectories of mice trained and tested in darkness versus trained and tested in light successfully traveling towards the target in both darkness (blue line) and light (red line) respectively.

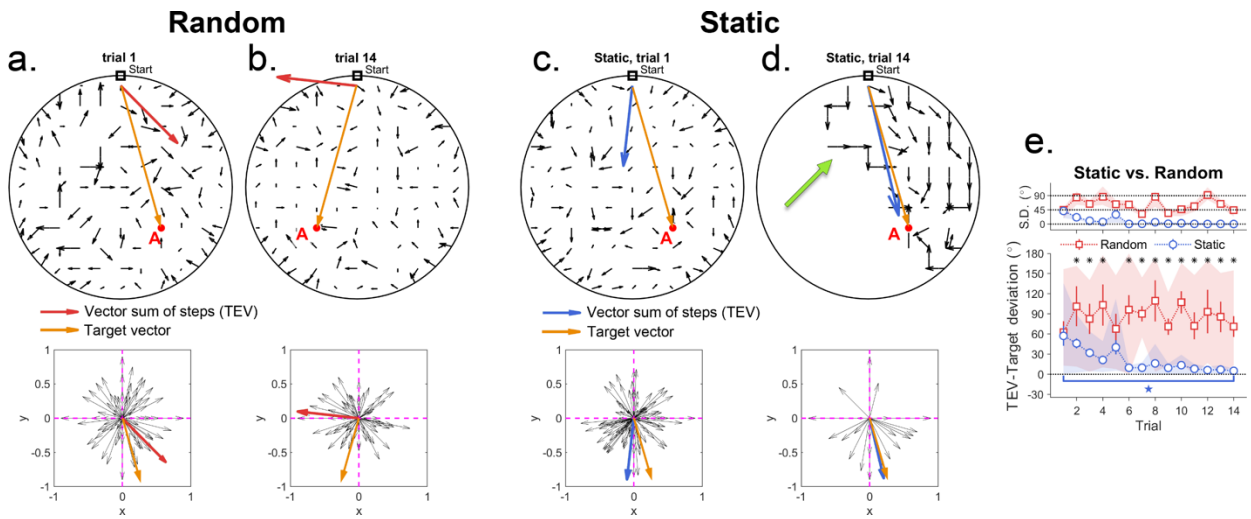

Extended Data Figure 4: Step map vectors in detail for the random and static entrance experiments. Arrows are the gradient of the step probability, and point to the most probably step starting from a given site in the arena. **a and b.** (top) Step probabilities in the random entrance experiment. Same data as in Figure 4 but with no background color, highlighting the step gradient vectors. The vector sum of steps (TEV) is shown in red (i.e., the vector sum of all the probability gradient step arrows), while the target vector is shown in orange (the vector pointing from the entrance to the direction of the target). The TEV reflects the learned direction and sense of the target. Its intensity is rescaled by the distance to the target for display purposes. (bottom) All the arrows from panel a starting from the origin (mouse entrance), highlighting that there is a preferential direction only for trial 14 of static entrance experiments (right panel). Orange arrows are the target vectors. In both early training (panel a) and late training (panel b), the TEV does not correspond with the food vector, highlighting a random search pattern. **c and d.** TEV (in blue) and food vector (in orange) during static entrance experiment. The TEV becomes closely aligned with the food vector during late training. The green arrow illustrates how trajectories first diverge from the TEV but then return to it. This pattern is approximately symmetric and so results in a TEV pointing almost directly to the food (error is 11 cm – the size of the lattice site, corresponding roughly to the average distance between nearest holes). **e.** Comparison of the difference in angle between the TEV and the food vector as training progresses. Angle differences decrease in the static case but not in the random case. N = 8 per group, significant differences are marked with an asterisk (FDR corrected, q<0.01). The star marks significant difference between first and last trials only for the static entrance case.

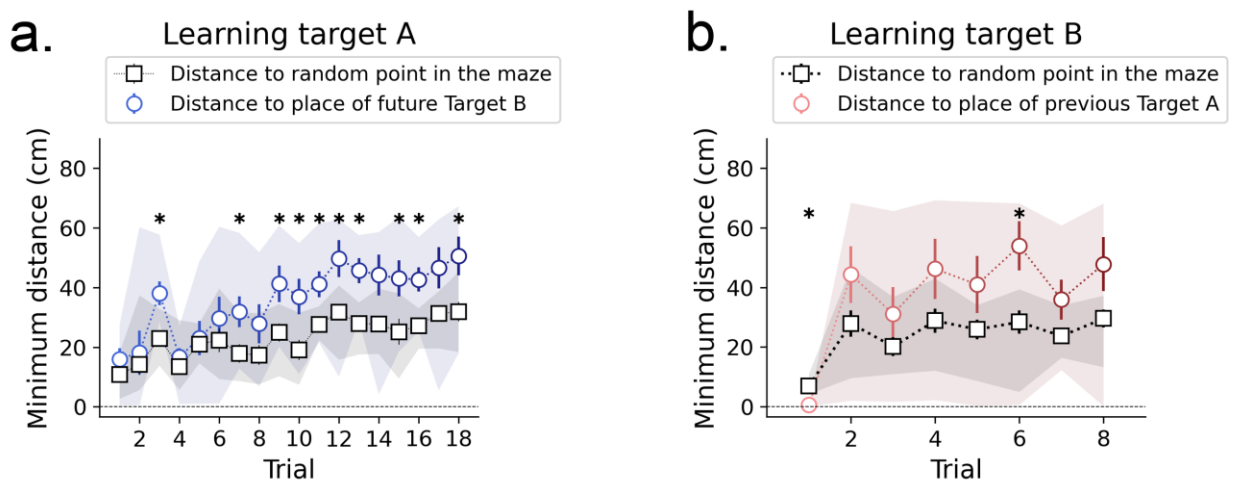

Extended data Figure 5: Minimum distance to alternative target in 2-target condition. a. During the learning of target A (no previous target has been presented), the mice keep consistently far from the location of where Target B will be placed. The minimum distance to the place where Target B will be placed is given by the blue circles; black squares are the minimum distance of the mice to a random point in the arena, averaged over ten random points, so it represents the minimum distance to a point expected by chance. **b.** The same pattern repeats when learning target B (target A has been already learned and is not present anymore): the mice keep consistently far from the location where the previous Target A was found. The minimum distance to the location of the previous Target A is given by the red circles; it is larger than the expected by chance, which is given by black squares and is calculated in the same way as in panel a. Shaded area represents the range of values, and bars represent the standard deviation. During the learning of B, some mice did visit the previous location sporadically, making the shaded area larger in panel b. However, they did not rely on target A to find B.

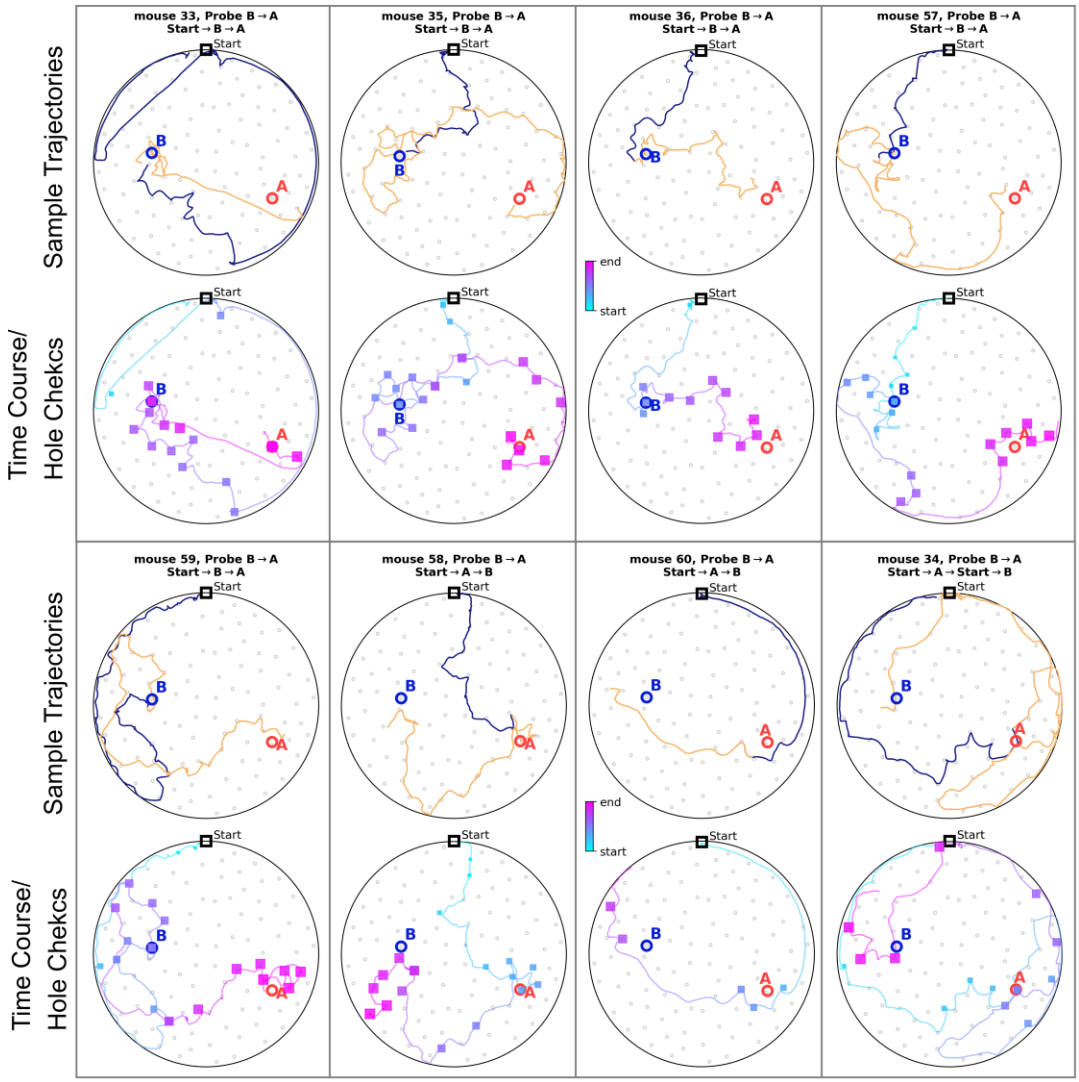

Extended data Figure 6: Trajectories and active sensing for the Probe B→A trial. Sample Trajectories: Each panel shows the mouse trajectory from the Start to 3 seconds after the moment where it gets the closest to either target A or B (whichever comes later, to ensure we observe its visit to both A and B locations). The orange trajectories are the ones that the mice took after realizing there is no food in either A or B (whichever comes earlier). We expect that mice perform the sequence Start→B→A, since B was trained last. Notice that mice 33, 35, 36, 57 and 59 do exactly this without ever going back to Start. Mice 34, 58 and 60 perform Start→A→B, but only mouse 34 visits the Start before heading from A to B. Time Course: Same as in Sample Trajectories panel but highlighting the time course (cyan: start; magenta: end) of the trial and the detected hole-check events (filled squares).

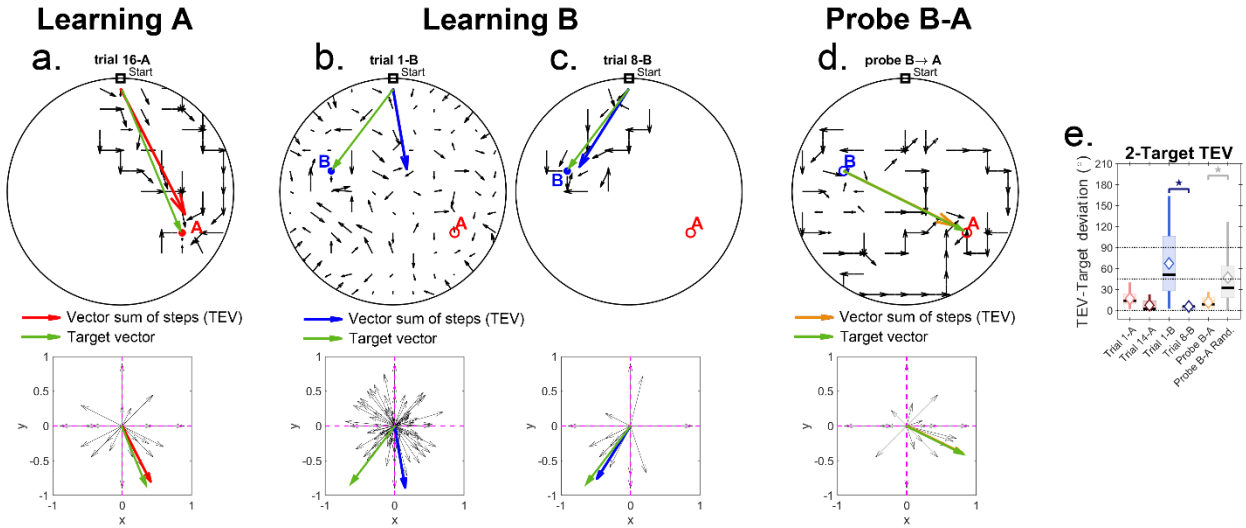

Extended data Figure 7: Step map vectors in detail for the 2-target experiments. Arrows are the gradient of the step probability, and point to the most probable step starting from a given site in the arena. **a.** (top) Same data as in Figure 6 but with no background color, highlighting the step gradient vectors. The vector sum of steps (TEV) is shown in red (i.e., the vector sum of all the probability gradient step arrows). The food vector (direction from entrance to food location) is shown in green. The TEV reflects the learned direction and sense of the target. Its intensity is rescaled by the distance to the target for display purposes. (bottom) All the arrows from panel a starting from the origin, highlighting that there is a preferential direction in trial 14-A (since this is analogous to static entrance experiment), 8-B and Probe B-A. Red arrows are the net displacement vector. **b and c.** Step probability gradient while the mouse is learning Target B. During the first trial, TEV (blue arrow) is still somewhat pointing towards target A. By late training, TEV closely corresponds with the food vector for target B. **d.** Step probabilities during Probe B-A. TEV (orange vector) points from target B back to target A, highlighting the short cut routes. **e.** The difference in angle between the food vector and TEV. The difference becomes significantly lower as training progresses, going back up as the target is switched, and then going down again as the mice learn the new target. During the probe trial, the TEV also significantly more closely matches the vector between target B and target A compared to a TEV generated using a random step map (gray box).

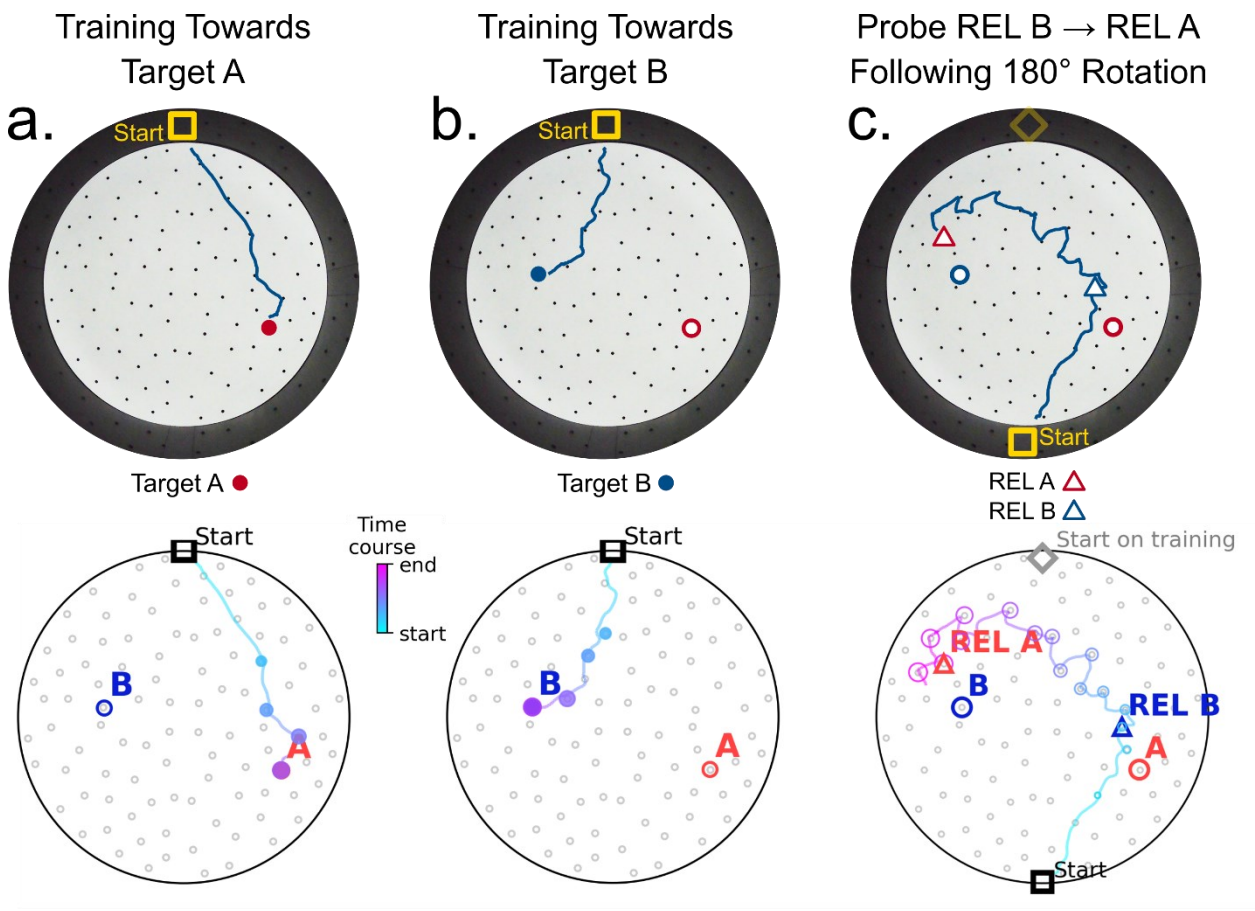

Extended data Figure 8. Two food location experiment plus 180° rotation. (**top**) Trajectories during training. (**bottom**) Hole checking corresponding to the above trajectory. a. The mouse's trajectory after successful training on Target A. b. The mouse's trajectory after food is removed from A and the mouse is successfully trained on Target B. c. The mouse's home is now rotated by 180°. The mouse first goes directly to the REL for Target B and subsequently goes directly to the REL for Target A, i.e., the combination of a two-food site protocol followed by an entrance rotation protocol results in trajectories that independently follow the rules for each protocol. Hole checking was also consistent with initial hole checking at Target A and the Start to A trajectory; After training on Target B, hole checking transferred to B and the Start to B trajectory. Finally, after the 180° rotation, hole checking occurred at the REL B and REL A sites and the trajectory connecting them.