Mice learn multi-step routes by memorizing subgoal locations

Philip Shamash[1], Sarah F. Olesen[1], Panagiota Iordanidou[1], Dario Campagner[1], Banerjee Nabhojit[1], Tiago Branco[1,*]

[1] UCL Sainsbury Wellcome Centre for Neural Circuits and Behaviour, London, W1T 4JG, UK

* For correspondence: t.branco@ucl.ac.uk

Animals must rapidly gather spatial information about new environments so that they can quickly reach food or safety even when direct paths are unavailable. The behavioral strategies used to implement multi-step routes to goals in naturalistic settings are unknown. Here we show that mice spontaneously learn a subgoal memory strategy while escaping to shelter or seeking food in an obstructed environment. We first investigated how mice navigate to shelter in response to threats when the direct path is blocked by a wall. Initially, mice ran straight toward the shelter and circumvented the obstacle using sensory cues. Over the course of 20 minutes, however, they switched to a spatial memory strategy to execute spatially efficient paths. Efficient escape routes were not learned by reinforcing egocentric actions or by constructing an unbiased internal map during exploration. Instead, mice used a hybrid strategy: they memorized specific subgoal locations encountered during previous running movements toward the obstacle. We then found that the same behavioral strategy is also used in a reward-seeking task. These results show that spontaneous memorization of local subgoals is a fundamental strategy by which rodents execute efficient multi-step routes to goals in novel environments.

Introduction

For prey species such as mice, quickly finding effective routes to goals is critical for survival because it reduces exposure to potential predators (Lima and Dill, 1990). This is a challenging task: natural environments are complex, and wild animals must compute multi-step routes taking into account uneven terrain, obstacles, and dynamically changing environments. Ethological studies of wild rodents have emphasized the roles of locating salient landmarks (Drickamer and Stuart, 1984; McMillan and Kaufman, 1995) and adhering to familiar paths (Benhamou, 1991; Thompson, 1982) in overcoming these challenges. However, observational studies are unable to identify the cues and behavioral strategies that animals actually use to navigate.

Experimental evidence from rodents trained to locate goals has uncovered multiple types of spatial reasoning that can be used to solve complex navigational problems. On the one hand, rodents can keep track of where they are within an allocentric (environment-centered) reference frame (Morris, 1981). This sense of position is thought to be

integrated into an internal topological map connecting places within the environment, which allows animals to compute subgoal locations whenever a new multi-step route toward a goal is required (Edvardsen et al., 2020; Spiers and Gilbert, 2015; Stachenfeld et al., 2017; Tolman, 1948). This kind of cognitive-map-based reasoning is flexible, is learned by observing the structure of the environment, and depends on the hippocampus (O'Keefe and Nadel, 1978). Alternatively, animals can navigate to goals without relying on an internal map. These strategies include integrating self-motion cues to compute a vector back to their starting position (Etienne and Jeffery, 2004); repeating egocentric movements at familiar junctions (Restle, 1957); and using landmarks for visual guidance (Hamilton et al., 2004). The latter two tactics, known as "taxon" strategies, are inflexible, rely on proximal cues, and are learned through previous motivated actions (O'Keefe and Nadel, 1978).

Despite all that is known about rodent navigation, the behavioral strategies that animals spontaneously use to quickly build up and deploy spatial knowledge in new environments remain unknown. The abilities listed above have mostly been demonstrated by repeatedly placing rodents in constrained mazes until they learn to navigate to a goal. In a natural setting, however, spatial learning must occur via internally generated exploration patterns and within a very limited timeframe. It is therefore unclear how well previous classifications of navigation strategies map onto the instincts and learning procedures that animals use during natural goal-directed navigation.

Escape behavior offers a powerful model for studying naturalistic navigation in the laboratory. Diverse animals, including fishes, lizards, crabs, birds, and rodents, respond to threats by escaping to a familiar shelter (Cooper Jr. and Blumstein, 2015). Mice are known to rapidly identify and memorize shelter locations in new environments and instinctively respond to visual or auditory threats by running straight to the shelter (Vale et al., 2017; Yilmaz and Meister, 2013). Previous studies have shown that the spatial memory for running back to shelter ('homing') can be based on path integration or distal visual landmarks when a direct path is available. (Alyan and Jander, 1994; Etienne et al., 1985; Harrison et al., 2006; Vale et al., 2017). If the direct path is blocked on one side by a barrier, previous work has shown that gerbils can use spatial memory to reach the hidden shelter after a brief period of exploration (Ellard and Eller, 2009). Thus, rodent escape offers not only reliable, stimulus-locked trajectories and rapid learning within a single session but also a reliance on spatial reasoning. These qualities make escape a useful model for understanding how animals learn and execute complex goal-directed trajectories within the time constraints compatible with survival in natural settings.

Here we first investigate the strategies that mice use to navigate to a shelter in response to auditory threats when the direct path is blocked by a wall. Through quantitative analysis of escape trajectories and their relationship to exploratory behavior, systematic variation of spatial conditions, and dynamic modifications to the environment, we describe how mice learn to execute efficient multi-step escape routes within minutes of entering a novel, obstacle-laden environment. We then show that the navigational strategy for escape is also used for reaching a food reward goal.

Results

Mice rapidly learn efficient escape trajectories in the presence of obstacles

As a baseline condition for investigating how mice learn escape trajectories, we placed naïve animals in a circular, open-field arena with a shelter. After a brief exploration period during which mice spontaneously located the shelter, we exposed them to a loud, overhead crashing sound while they were in a pre-defined threat zone (Fig. 1A). This reliably elicited rapid escapes directed at the shelter along a straight 'homing vector' (Fig. 1A-B, Extended Data Fig. 1A, Supplementary Video 1-2), similar to previous results (Vale et al., 2017). We then repeated this experiment in a separate group of mice, with a wall positioned between the threat zone and the shelter. This wall was white against a black background and all mice approached and walked along it during the exploration period (Extended Data Fig. 1B). To quantify initial escape directions, we computed a direction score between 0 and 1, where 0 is an escape vector aimed directly at the shelter and 1 is a vector aimed at the obstacle edge (Fig 1A); escapes are classified as edge vectors if they surpass the 95th percentile of escape direction in the open field (direction score > 0.68). On the first trial of threat presentation, 53% of the mice ran toward the shelter as in the open-field condition and only circumnavigated the obstacle after approaching it ('homing-vector escapes; Fig. 1A-B; Supplementary Video 3; median direction for all mice: 0.56, IQR 0.25 - 0.92, N=17 mice), while 47% escaped by running directly to the edge of the obstacle and then to the shelter ('edge-vector escapes; Supplementary Video 3). Differences in escape direction were not related to differences in position or angle at the moment of threat onset (mean distance from center of wall: 38.7 cm for homing vector vs 38.9 cm for edge vectors, P=0.84, *permutation test*; mean body angle relative to shelter: 105º for homing vectors vs 110º for edge vectors, P=0.76). Despite having previously explored the obstacle, the escape trajectories of mice with homing-vector responses resembled escapes when an unexpected obstacle arises at the moment of threat onset (Extended Data Fig. 1C, Supplementary Video 4). This suggests that mice that initiate escapes following the homing vector rely on immediate sensory input to negotiate the barrier. Next, we performed two control experiments to test whether homing-vector escapes were directed at the shelter or were instead exclusively aimed at the safety that can be provided by running along a wall (Simon et al., 1994). First, replacing the wall obstacle with an unprotective hole obstacle did not reduce the fraction of homing-vector responses (Extended Data Fig. 1D-E). Second, moving the shelter to the side of the arena abolished escapes directed at the center of the obstacle (Extended Data Fig. 1D-E). Thus, after limited experience with an obstacle, a large proportion of mice default to homing-vector responses when exposed to threats.

Over the course of 3 threat presentation trials within 20 minutes, mice performed escapes that were increasingly spatially efficient (median ratio of shortest path to the actual escape path; trial 1: 0.77, IQR 0.65-0.89 vs.

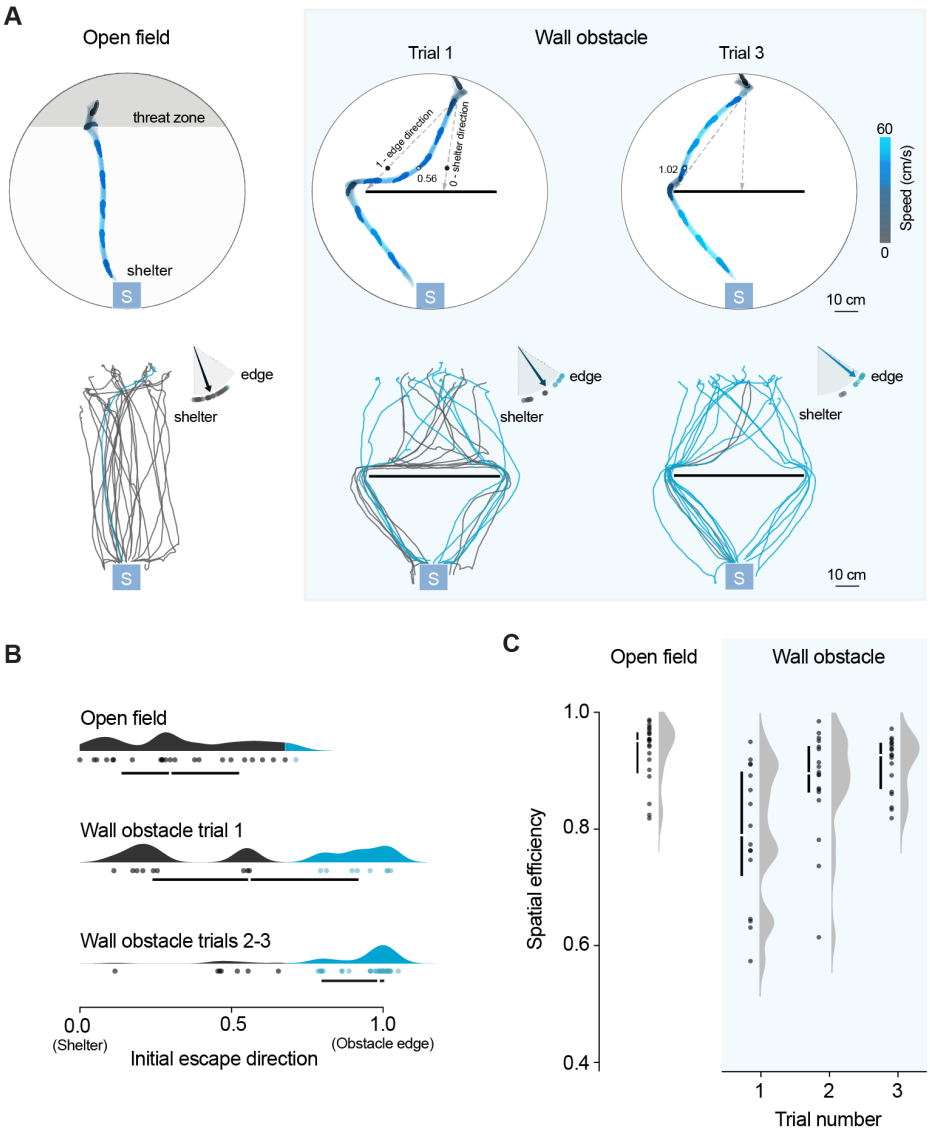

Figure 1 - Mice rapidly learn efficient escape trajectories in the presence of obstacles

(A) Single escape trials color-coded by speed (top) and all trajectories from each condition color-coded by trajectory type (bottom; gray: homing-vector paths, blue: edge-vector paths). Initial escape direction is the direction of the vector between escape initiation and 10 cm in front of the obstacle location, normalized between 0 (directed at the shelter) and 1 (directed at the obstacle edge location). Dashed arrows in top panels (middle and right) illustrate shelter and edge vectors directions, and dots show the vector end-point for direction measurement. Values (0.56 and 1.02) are the escape direction score for the examples shown (values >1 indicate over-shooting the edge location). Arrow plots show the distribution of initial escape directions (arrow is the median value). Obstacle: N=19 mice, 53 trials; open field: N=10 mice, 23 trials. **(B)** Summary plots of initial escape direction. Mice execute homing vectors in the open field and upon first experiencing an obstacle. They learn to execute edge-vector escapes past an obstacle within three trials (trial 3 is 17±4 minutes into the session, mean±std). Direction score on trial 1 vs. trials 2-3: P=0.009, permutation test. Each dot represents one escape (2-3 escapes per mouse). Black points are escapes classified as homing vectors and blue points are edge-vectors escapes. Lines show the median and IQR. Distributions are kernel density estimates. **(C)** Escape spatial efficiency computed as the ratio of the shortest path to the shelter over the path actually taken to the shelter. Efficiency is initially lower with an obstacle, but increases over the course of three trials. Each dot represents one escape. Lines show the median and IQR. Distributions are kernel density estimates.

trial 3: 0.92, IQR 0.86-0.95; F(2, 30) = 7.1, p = .003, *repeated measures ANOVA on trials 1-3;* Fig. 1C) and rapid - (median normalized escape duration: 3.6 s, IQR 3.1-5.1 s on trial 1 vs 2.7 s, IQR 2.5-3.0 on trial 3; F(2, 28) = 8.3, p = .001; Extended Data Fig. 1F). By this point, almost all trajectories were aimed directly at the obstacle edge (median direction: 0.96, IQR 0.81 – 1.00), generating escape routes as efficient as those in the open-field arena (P=0.24 for trial 3 vs. open field, *permutation test*). Thus, while instinctive responses initially generate a large fraction of slow, inefficient escapes, limited experience with the environment is sufficient for naïve mice to acquire rapid and spatially efficient routes to the shelter.

Mice use a spatial memory strategy to learn efficient obstacle avoidance

We next investigated whether the increase in performance requires visual input to locate and run toward the obstacle edge. Presenting threats in complete darkness to naïve mice resulted in homing-vector responses similar to the open field condition, even if the initial exploration period with the obstacle was done in the light (% homing-vector escapes in the dark: 71% in the open field, 71% with the obstacle, 72% with the obstacle and exploration in the light; Extended Data Fig. 2A; Supplementary Video 5). This indicates that visual input is required during escape in order to efficiently avoid an obstacle when experience is limited. However, mice with a prior session with an obstacle performed edge-directed responses above chance levels despite the lack of visual cues (53% edge-vector escapes, P=0.025, *permutation test on direction*; Extended Data Fig. 2B-C). We thus considered that learning efficient escapes might entail developing a memory of the obstacle edge location, rendering visual perception of the obstacle unnecessary. To test this hypothesis, after the animals explored the environment with the obstacle for 20 minutes and with 3 escape trials, we removed the obstacle at the moment of threat onset ("acute obstacle removal"; Supplementary Video 6). Despite the obstacle disappearing before the termination of the mouse's initial orientation movement, all animals escaped along the edge vector and did not turn toward the shelter until they passed by the location where the obstacle edge used to be (median direction = 0.98, IQR 0.84-1.01; N = 7 animals; Fig. 2A-B). This experiment shows that after experience in the arena, edge-vector escapes have a strong memory component. Next, we examined how persistent this memory-based strategy is. We allowed mice to explore for ~5 minutes after removing the obstacle ("chronic obstacle removal"; Supplementary Video 6), during which time they all visited the now empty center of the arena (Extended Data Fig. 3A). Escapes on 46% of the subsequent trials were still directed at the location where the obstacle edge used to be (N = 29 trials, 13 animals, 2-3 trials per animal; Fig. 2A-B), indicating that memory-based escape trajectories persisted even after exploring the arena with the obstacle removed.

Developing a memory for guiding edge-vector escape trajectories could in principle be acquired through repeated escapes during exposure to threat or through exploratory behavior alone. To distinguish between these two possibilities, we let mice explore the arena with the obstacle for 20 minutes without exposing them to threat stimuli.

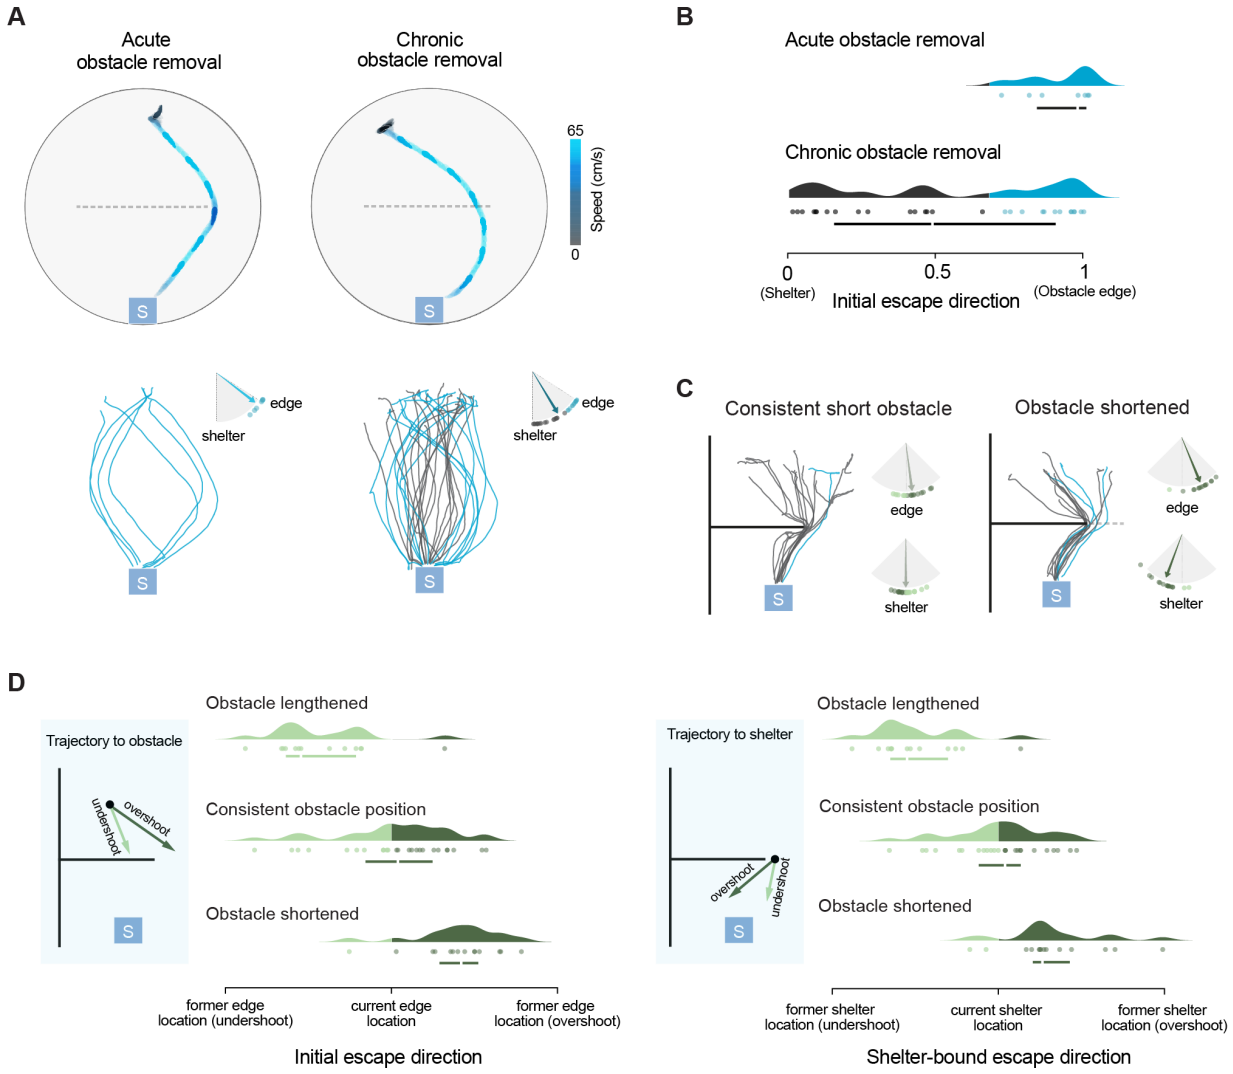

Figure 2 - Mice use a spatial memory strategy for efficient obstacle avoidance

(A) Example escape trials (top) and all trajectories (bottom) after removing the obstacle. Mice execute escapes directed at the obstacle edge in all trials when the obstacle is removed acutely at the threat onset, and in half of the trials after having several minutes to explore the newly open-field arena. P=0.004 for chronic obstacle removal vs. open field, permutation test on probability of edge-vector responses. Dotted line indicates where the obstacle used to be. **(B)** Summary plot for initial escape direction after obstacle removal. **(C)** Escape trajectories and direction after decreasing the length of the obstacle (right) and for a constant obstacle length (left). Escape paths consistently overshoot both the obstacle edge and the shelter when the obstacle is shortened. Grey indicates trajectories targeting the shorter obstacle edge, and blue indicates trajectories targeting the longer obstacle edge. In the arrow plot, dark green indicates overshooting the current edge or shelter position, and light green indicates undershooting. **(D)** Summary data for the trajectories shown in **C** and for a complementary experiment in which the obstacle gets longer. If mice start with a long obstacle that becomes shorter, escapes overshoot the obstacle edge (P=0.002, permutation test on direction) and the shelter (P=0.001). Alternatively, if mice start with a short obstacle that becomes longer, escapes undershoot the obstacle edge (P=0.002) and the shelter (P=0.0002).

Threat presentation after this period resulted in persistent edge-vector responses that were as frequent as when threats were presented during exploration with the obstacle (Extended Data Fig. 3B). These results show that within 20 minutes in a novel environment, mice spontaneously develop a robust strategy for multi-step escapes that relies on spatial memory.

Our experiments so far have revealed a memory-based obstacle avoidance strategy. We have also shown, however, that sensory input is used to navigate around the obstacle when experience is limited. We therefore tested how perception and spatial memory operate in tandem when both are fully available. We repeated the chronic obstacle removal experiment in a novel arena, but instead of removing the obstacle, we changed its length by 25% (Fig. 2C-D, Extended Data Fig. 3C; Supplementary Video 7). As expected from the previous results, initial escape trajectories were biased toward the previous edge locations (median direction after length change = 0.44, IQR 0.26-0.61, where 0 is previous edge location and 1 is the new location; N=27 trials, 19 animals). Nevertheless, the behavior differed from the obstacle removal experiments in two ways. First, the memory bias was intermediate in magnitude; unlike with obstacle removal, escape directions are not shifted all the way to the former edge location (Fig. 2D). Second, the second step of the escape was also biased by experience. After reaching the new edge location, mice ran in the direction of where the shelter would have been relative to the edge if the edge had not moved (median direction after length change = 0.37, IQR 0.25-0.67, where 0 is previous shelter location and 1 is the new location; Fig. 2D). These results show that, when available, the current obstacle position is an important cue for anchoring and adjusting memory-guided paths to both the edge and the shelter.

Mice learn to avoid obstacles during escape by memorizing subgoal target locations

We next aimed to characterize the spatial-memory strategy and how it is learned. We evaluated whether mice were learning escape routes using three possible strategies: habitual learning of turn angles, constructing a global cognitive map of the arena, or learning subgoal locations individually. First, we tested whether mice learn egocentric movements from the threat zone to the obstacle edge, similar to habitual learning in mazes. To test this, we extracted from the chronic obstacle removal experiment all continuous turn-and-run movements from the threat zone to within 5 cm of the obstacle ("homing movements"; Fig. 3A,D, Extended Data Fig. 4A; Supplementary Video 8) . These included both stimulus-evoked escape trajectories and similar movements performed spontaneously during exploration. After validating a linear prediction model to predict escape direction from turn angles (Extended Data Fig. 5A-B), we tested whether mice could be implementing escapes by repeating turn angles from previous homing movements. We took the previous turn angle that is most similar to the one used during the escape and used it to predict escape direction, i.e. whether the mouse will escape toward the shelter or toward an obstacle edge. We found that these angles did not match escape angles precisely enough to reliably predict escape direction (median R-square from linear prediction model = -0.03; Fig. 3B). Thus, homing movements are not stereotyped enough to explain escape directionality, and therefore reinforcement of turns cannot explain memory-guided edge-vector responses.

To experimentally verify that memory-guided edge-vector responses are not habitually repeated responses but rather goal-directed actions, we examined whether their expression is sensitive to changes in the goal location. We

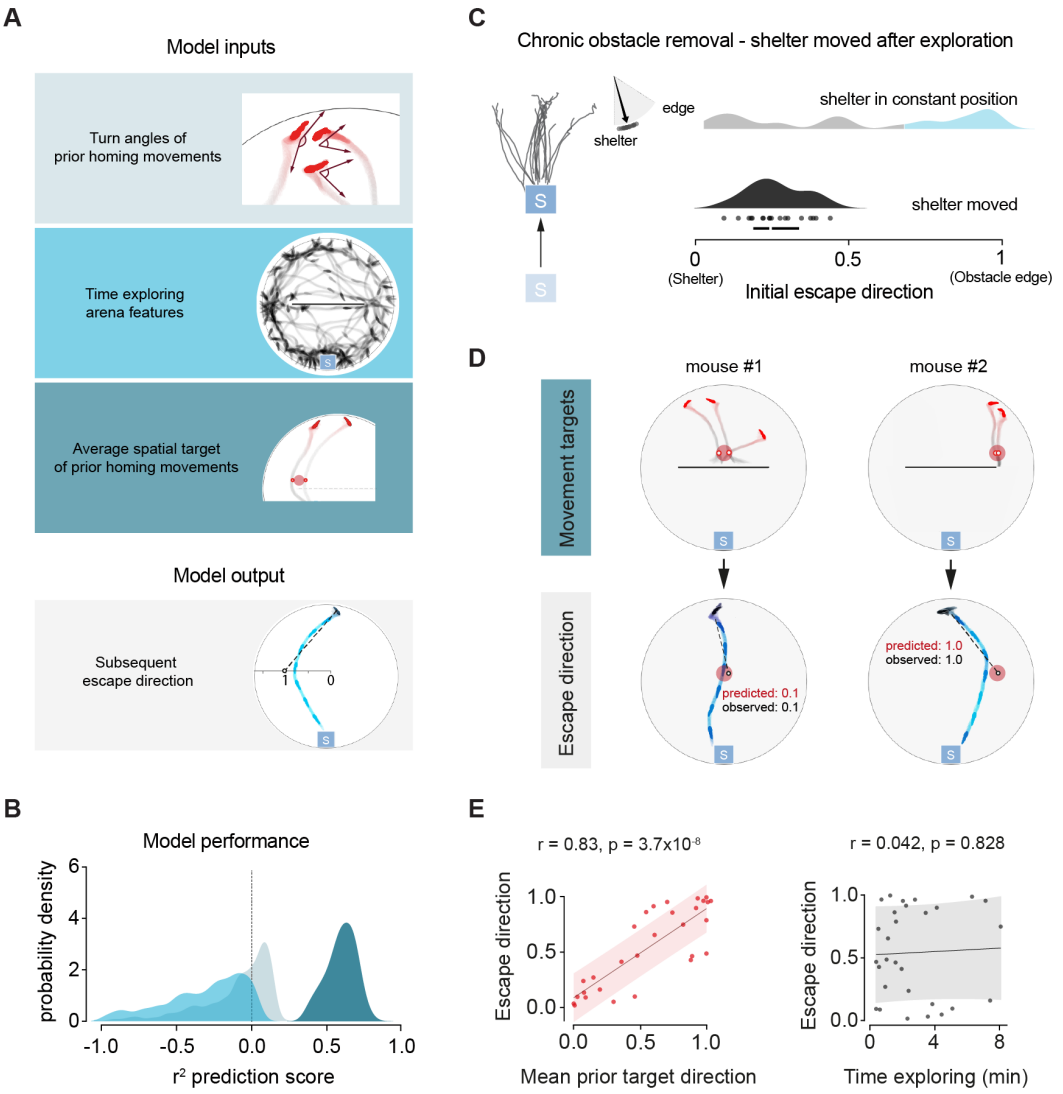

Figure 3 - Mice memorize sub-goal target locations rather than habitual actions

(A) Illustration of three inputs to a linear model that predicts escape direction after obstacle removal. Top: the angles of previous turn-and-run movements that best matches the subsequent escape angle is used for predicting escape direction. Dark red arrows are examples of angles of previous homing movements. Middle: time spent exploring features of the environment as a predictor. Image shows example exploration trajectories for a mouse in the presence of a shelter and an obstacle. Bottom: prediction using target locations of previous homing movements (small red circles). The large red circle indicates the mean target location. Model output: escape direction prediction (dotted line and small black circle; 0 is escape in shelter direction, 1 is escape towards obstacle edge). **(B)** Distribution of prediction scores for the different inputs to the model. Escape direction can be well predicted from previous target locations but not from previous turning movements or from time spent exploring features of the arena. **(C)** Escape trajectories and summary data for escape after removing the obstacle and simultaneously changing the shelter location. Escapes are no longer persistently directed at the obstacle edge. **(D)** Homing movement history for two different mice from the chronic obstacle removal experiment, and subsequent escape trajectories. Mean prior target direction (large red circle) is computed as the average of previous prior homing-movement target locations. This location is used to predict the escape direction. The small black circle and dotted line indicate the actual direction of the observed escape. **(E)** The direction predicted by prior movement targets correlates well with subsequent escape directions, but time exploring after obstacle removal does not. Shaded area shows the prediction interval within 1 standard deviation of the mean.

repeated the chronic obstacle removal experiment, but this time moving the shelter to a new location after the 20-minute period with the obstacle. The distribution of egocentric actions during the initial 20 minutes was the same between the constant shelter and shelter displacement conditions (Extended Data Fig. 4B), which should result in a similar frequency of edge-vector responses if the underlying process were based purely on action reinforcement. In contrast, all escape paths after moving the shelter were homing vectors directed toward the new shelter location (median direction = 0.25, IQR 0.20-0.34, N = 7 mice, 18 trials), with zero edge-vector responses (Fig. 3C). This experiment confirms that the spatial memory mechanism for escape does not reinforce habitual, egocentric movements. Instead, our data suggest that edge-vector responses result from a goal-directed process, where mice specifically target the obstacle edge location in order to reach the shelter.

Next, we investigated whether learning the spatial relationship between the obstacle edge and shelter locations arises from constructing an abstract internal map of the arena during the exploration period. Thus, we measured the correlation between escape paths and the time spent exploring various features of the environment, such as the obstacle, the obstacle edge, the back half of the arena and the entire arena, both before and after obstacle removal. Exploring the environment should help incorporate its features into an internal model of space, and so these correlations could provide evidence that mice use such an abstract representation for escape. However, none of the 12 correlations analyzed were significant (median correlation coefficient 0.06, IQR 0.00–0.19; median p-value 0.57, IQR 0.32–0.84; Fig 3E; Table 1).

We therefore examined the alternative possibility that instead of using a global map, mice learn to repeatedly target individual subgoal locations accessed during previous homing movements. To test this hypothesis, we extracted the target, rather than the angle, of homing movements (Fig. 3A,D). We first found that the occurrence of a prior homing movement targeting the obstacle edge was significantly correlated with the escape direction (R=0.52, P=0.004; Table 2). Next, we computed the mean homing movement target prior to the escape and used this target point to predict subsequent initial escape directions; this measures how often the mouse has previously run along homing or edge-vector paths (excluding edge-vectors to the obstacle edge on the opposite side of the arena; Fig. 3D). This prior movement target direction was even more strongly correlated with the direction of subsequent escapes (R=0.83, $P=3.7x10^{-8}$; Fig. 3E, Table 2). These correlations were only present when examining the relationship between escapes directed toward a particular obstacle edge (i.e. left vs. right) and prior homing movements toward that same side, but not with prior homing movements toward the opposite edge (Table 2).

To directly compare the predictive power of time spent exploring against average prior movement targets, we built a linear model to predict escape directions from exploration inputs and prior homing movements. For the exploration inputs, we used the best correlate with escape direction (time spent exploring the side of the arena opposite from the shelter), both before and after obstacle removal. The model showed that prior homing movements are highly predictive of the spatial targets of escape (median R-square = 0.60), while time exploring has no predictive power

(median R-square = -0.32; Fig. 3B, Extended Data Fig. 5C). Weighting prior homing movements more strongly if their initial position or angle was close to the escape's initial conditions did not improve prediction (Extended Data Fig. 5D). This affirms that mice are learning subgoal locations rather than actions or paths.

Finally, we verified these results on the acute obstacle removal experiment by removing the obstacle during the 1st trial. As expected by our analysis of the chronic removal experiment, all mice with at least one prior movement targeting the obstacle edge also ran to that edge location during the escape (median direction = 0.98, IQR 0.91-0.98; Extended Data Fig. 4C; Table 1). All mice with no such movements, however, ran straight toward the shelter (median direction = 0.30, IQR 0.29-0.43; Extended Data Fig. 4C). Additionally, none of the correlations with exploration were significant (Table 1). Together, these results suggest that the spatial-memory strategy for escape consists of directly learning specific spatial subgoal locations targeted by previous homing movements, rather than repeating fine-grained actions or gradually constructing an unbiased internal map through exploration.

Instinctive exploratory movements and one-shot learning can explain subgoal learning

The ability of our model to predict escape targets based on the prior history of edge-vector movements suggests that these movements might be the basis for rapid learning of spatial subgoals. However, it remains unclear what prompts these movements in the first place and how they may lead to the acquisition of subgoals. Thus, we further analyzed the properties of spontaneous edge-vector movements and their relationship to learned escape paths. First, we observed that edge-vector movements occur most during the first few minutes of the session and occur equally with or without a shelter in the environment (Fig. 4A-B). Second, when the obstacle is a hole, learning to perform efficient multi-step escapes takes twice as long (% edge-vectors escapes: 20% on trial 2-3, 67% on trial 6-7; Fig. 4C) and edge-vector movements occur with the same, lower frequency as in the open-field arena (Fig. 4B). Edge-vector movements thus initially consist of instinctive, exploratory actions directed at a salient edge in the environment. We next examined memory-guided escape directions after having executed different numbers of these movements. After obstacle removal, edge-vector escapes frequently occurred after 1 or more prior edge-vector movements from the threat zone (62% after 1 prior movement, median direction = 0.82, IQR 0.43-0.98) but never after zero (median direction = 0.24, IQR 0.11-0.29; Fig. 4D), suggesting that a single exploratory movement toward an obstacle edge can establish a subgoal for use in subsequent escapes.

Next, we analyzed whether exploratory movements toward or away from the shelter have a differential role in establishing subgoals. While mice execute edge-directed movements from both the shelter side and the opposite side of the arena, we found that movements from the threat area toward the shelter are much more predictive of subsequent escapes (toward shelter: $R = 0.83$, $P = 3.7 \times 10^{-8}$; away from shelter: $R = 0.43$, $P = 0.019$; Table 1). In addition, weighting the importance of prior homing movement targets based on how soon they arrive at the shelter improved the prediction

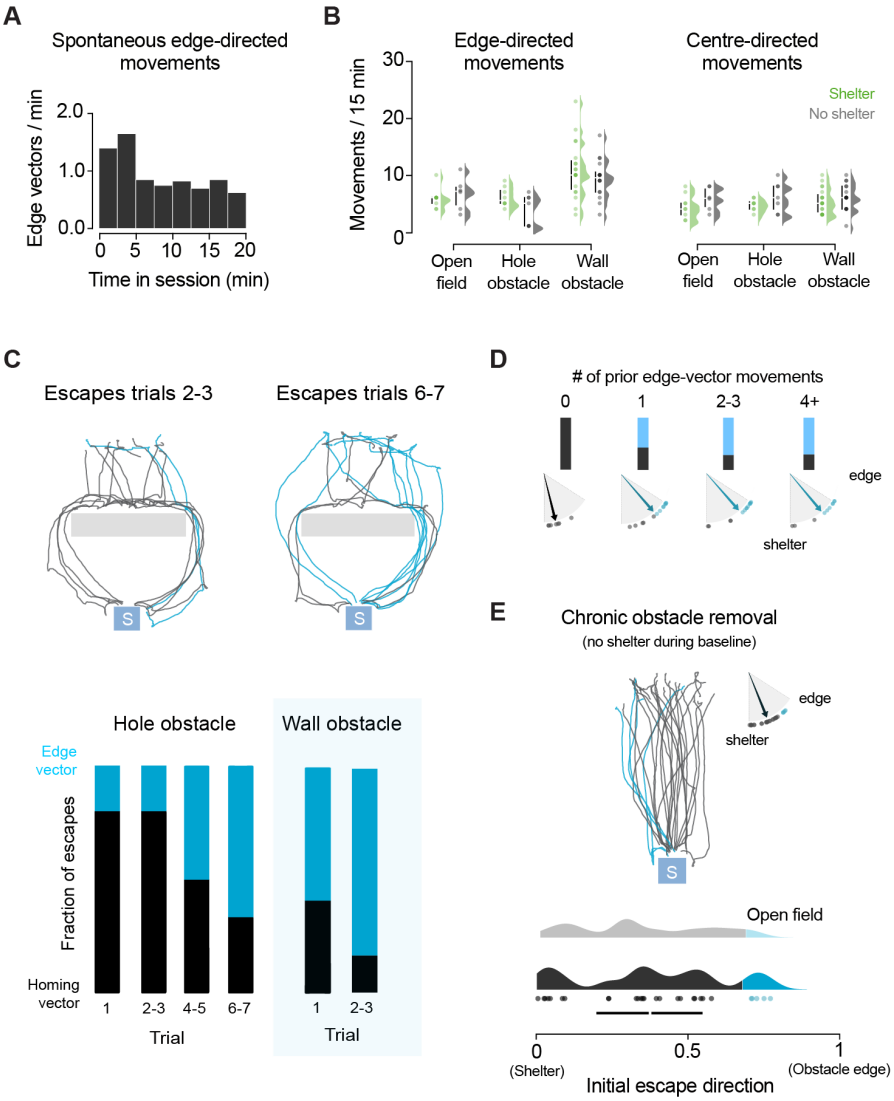

Figure 4 - Spontaneous edge-vector movements underlie sub-goal learning

(A) Frequency of spontaneous movements toward the obstacle edges during behavioural sessions (data is pooled from all sessions). **(B)** Frequency of edge-directed movements for different conditions. Edge-vector movements are enriched when wall-obstacle edges are present, including when there is no shelter in the arena (wall obstacle vs. open field: P=0.0005, permutation test), but not in the arena with the hole obstacle (hole obstacle vs. open field: P=0.90). For movements directed toward the center of the arena, there are no significant differences across conditions. Each dot is one session. Lines show the median and IQR. Distributions are kernel density estimates. **(C)** Escape trajectories for a hole obstacle (top) and evolution of escape trajectory type for increasing number of trials (bottom; wall obstacle data shown for comparison). Learning to execute efficient escapes takes longer when blocked by a hole obstacle rather than a wall obstacle. **(D)** Initial escape direction for increasing number of prior edge-vector movements across all obstacle removal experiments. Memory-guided escapes toward the edge location occurred after 1 or more prior edge-vector movements, but never after zero. **(E)** Escape trajectories and summary data for escapes with no obstacle and a shelter, after an initial exploration period with an obstacle but no shelter. Unlike in the original chronic obstacle removal experiment (see Figure 2A-B), the distribution of escapes is similar to the open-field condition (see Figure 1A-B).

of memory-guided escape direction (median R-square without weighting = 0.64; median R-square with weighting = 0.74; Extended Data Fig. 5D). These analyses suggest that subgoals are identified in a direction-specific manner and are based on providing access to a goal on the opposite side of the arena. To test this hypothesis experimentally, we first let mice explore the arena with a wall obstacle but without a shelter. After exploration, we removed the obstacle and added the shelter, which all mice quickly located. In contrast to the behavior when the shelter is present throughout the session, threat presentation under these conditions elicited escapes that were predominantly homing-vector responses (82%), with only a few edge-vector trajectories (18%; median direction: 0.38, IQR 0.20-0.55; Fig. 4E). The subgoal memory was therefore deployed infrequently. Nonetheless, escape direction was still significantly correlated with the mean prior movement target but not with any other exploration or movement metric (Tables 1-2). Thus, as predicted by our analysis of edge-vector movements, subgoal memorization occurs primarily when the shelter-goal is present during exploration. These results support a role for instinctive edge-vector movements together with knowledge of a goal location in generating the rapid learning of multi-step escape routes.

Our data show that mice learn efficient escape routes through a subgoal strategy that relies on execution of edge-vector movements toward the shelter. We next further characterized how spatial efficiency changes with experience. We found that the improvement in spatial efficiency with experience is fast, as a single prior edge-vector homing movement is significantly associated with greater escape efficiency in the presence of an obstacle (median efficiency with 0 prior edge vectors: 0.81, IQR 0.74-0.89; with 1 prior edge vector: 0.91, IQR 0.85-0.93, P = 0.002, *permutation test*; Figure 5A), without changing measures of motivation to escape (escape initiation time and speed) or other confounding variables (time exploring and trial number; Extended Data Fig. 6A-B). On the other hand, subgoal memorization causes mice to persistently take roundabout routes even after a new, direct path becomes available. Escape efficiency after obstacle removal is lower in mice with one prior edge-vector movement compared to those with none or to naïve mice (median efficiency with 0 prior edge vectors: 0.96, IQR 0.94-0.97; 1 prior edge vectors: 0.89, IQR 0.83-0.92, 0 vs. 1 edge vector: P = 0.008, *permutation test*; 1 edge vector vs. naïve mice in the open field: P=0.04; Figure 5A), again without changing confounding variables (Extended Data Fig. 6A-B). A cost of the rapid subgoal memorization strategy is therefore reduced flexibility to sudden changes in the environment.

Subgoal learning supports food-seeking trajectories

Based on this trade-off favoring speed over flexibility, we considered that this strategy could be specific to stimulus-evoked escape behavior. To test this, we performed an obstacle removal experiment in the context of a less urgent, reward-based task. First, we trained food-deprived mice to approach and lick a reward port in response to a 10-kHz tone, which indicated the availability of condensed milk at the port. This took place across 5 sessions, in an operant conditioning box (Extended Data Fig. 7). Next, we transported this task to the arenas previously used for escape

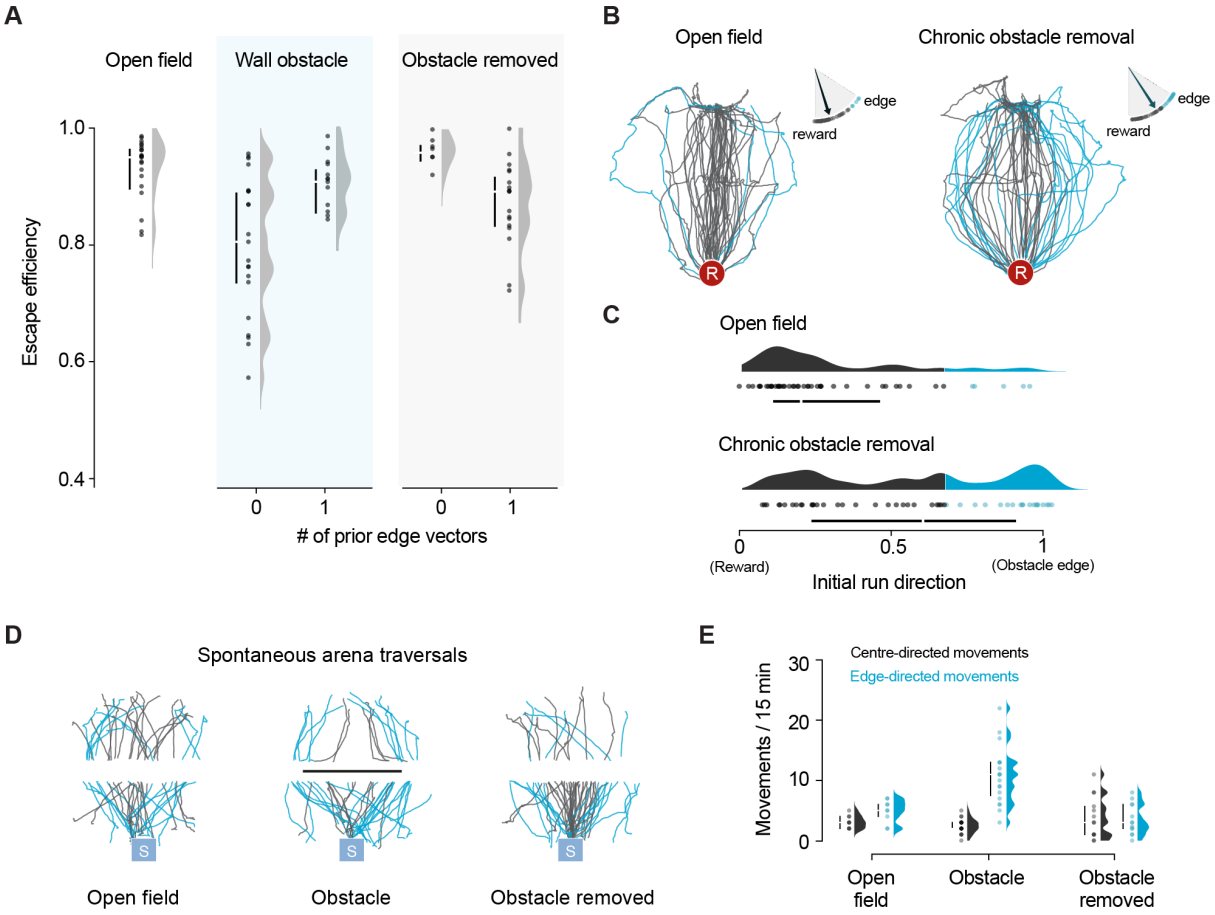

Figure 5 - Sub-goal learning supports efficient escape and food-seeking trajectories

(A) Escape efficiency, comparing escapes with zero vs. one prior edge-vector movements (toward the edge used during the escape). Escapes in the presence of an obstacle have higher efficiency after one prior edge vector. After removing the obstacle, having one prior edge vector is instead associated with lower efficiency. **(B)** Food-approach trajectories in response to a 10-kHz tone associated with the availability of condensed milk at the lick spout. The red circle with 'R' shows the reward location. **(C)** Summary data for food-seeking trajectories shows persistent edge-directed food-approach paths after experience with the obstacle during exploration. **(D)** Paths taken during spontaneous exploration in the escape experiments, from the ends of the arena toward the center (conditions with more sessions are randomly downsampled so that the same number of paths is displayed for each condition). **(E)** Number of spontaneous center-directed and edge-directed movements during exploration. Edge-directed movements in the chronic obstacle removal experiment are not increased compared to the open-field condition. Open field vs. obstacle: P=4x10-5; open field vs. obstacle removed: P=0.68, permutation test on number of edge-directed movements.

behavior: the shelter was replaced by the reward port, and the threat stimulus was replaced by the 10-kHz tone. Mice reliably ran toward the reward port upon tone presentation, but with slower reaction times (median time to start running toward goal: 1.5 s, IQR 0.7-3.5 s for food task; 0.6 s, IQR 0.4-1.2 s for threat response; P=0.005, *permutation test*). Similar to escape routes, following 20 minutes of experience with the obstacle and its subsequent removal, a large proportion of food-approach paths were edge-vector responses (40% edge vectors; median direction: 0.61, IQR 0.24-0.91; P = 0.01 compared to open field, *permutation test on probability of edge-vector responses*; Fig. 4A;

Supplementary Video 9). As with escape, these trajectories were correlated to prior running movement targets but not with other features of exploration (Tables 1-2).

Finally, we tested whether experience with the obstacle induces a non-specific increase in edge-vector movements, as this could trivially explain the apparent use of subgoal memorization across two distinct tasks. We measured how frequently mice moved from the ends of the arena toward the center and toward the edge locations during exploration. In contrast with evoked escapes and food-approach paths, exploratory paths following obstacle removal were no more likely to target the obstacle edges than in the open-field control condition (Fig. 5D-E). Subgoal memorization therefore reflects a strategy for goal-directed navigation rather than a general bias in how mice move around their environment following experience with an obstacle. These experiments suggest that subgoal memorization and its underlying trade-off between speed and flexibility are general components of how mice implement multi-step paths to goals in novel environments.

Discussion

We have found that mice learn efficient routes to goals in obstacle-laden environments using a fast, innate subgoal memorization strategy. For the first few minutes after entering a new environment with obstacles, mice escape to shelter by relying on their homing instinct and memory of the shelter location to run directly toward the shelter, and then use sensory cues to navigate around obstacles as they are encountered. These escapes are spatially inefficient. They resemble paths that species with less advanced spatial reasoning, such as toads, crabs, and ant colonies, take around obstacles to reach a goal (Collett, 1982; Freas and Schultheiss, 2018; Layne, 2003). Over the course of 20 minutes, however, mice start taking escape routes that directly target the obstacle edge. These efficient escapes depend on a memory of the edge location; this is particularly evident when mice complete multi-step escapes around an obstacle even after the obstacle has been removed.

Previous work has shown that rodents use spatial memory to navigate to shelter in an open-field arena (Alyan and Jander, 1994; Etienne et al., 1985; Vale et al., 2017). In such a simple environment, however, escape routes can be implemented by path integrating self-motion cues to keep track of a single vector to the shelter location – a purely egocentric process. With obstacles in the environment, a more advanced strategy is needed. Previous results in gerbils escaping in an obstructed environment suggested that spatial memory was employed to reach the shelter (Ellard and Eller, 2009), but the navigational strategy that the animals used was unknown. Our results show that mice efficiently navigate to shelter by using subgoals in an allocentric reference frame. Several observations support this view. First, mice accurately target the edge location at least 10 minutes after the obstacle has been removed, which cannot be explained by a pure egocentric strategy. Second, escape trajectories involved immediately orienting and running toward

a sub-goal ~50 cm away, which is not consistent with following odor trails or gradients (Liu et al., 2020; Wallace et al., 2002). Third, mice switch between targeting the shelter and the subgoal depending on which side of the obstacle they are on. This requires awareness of external metrics such as vectors to the obstacle and shelter. Fourth, the inability of turning movements to predict memory-guided escape directions rules out repeating egocentric turns as a primary mechanism. Fifth, moving the shelter after the exploration period abolished the expression of memorized subgoals, showing that the spatial memory strategy is goal-directed rather than habitual. Finally, when we changed the obstacle's length, escape paths inaccurately targeted not only the new obstacle edge location but also the shelter. This suggests that mice are tracking the spatial relationship between the obstacle edge and the shelter within their spatial schema of the arena. Future experiments on how efficient escape routes transfer across different obstacles and tasks may further specify the nature of this allocentric schema and how it is updated.

Experimental work on rodent navigation typically frames allocentric navigation as a process of incorporating perceptual landmarks into a map, from which novel routes can be computed. There are, however, heuristic alternatives to unbiased map-based reasoning, which can be revealed by analyzing trajectories and prior exploratory patterns (Janson, 2014; Teichroeb and Smeltzer, 2018). Rodent navigation studies have generally not examined the relationship between initial exploration and subsequent navigation strategy in naïve subjects. While subgoal learning does rely on an allocentric spatial schema, our analysis of exploration did not support the view that mice compute escape routes by gradually building up an unbiased representation of the arena or an 'obstacle to be avoided' feature. In particular, none of the following variables led to an obstacle edge being tagged as a subgoal: 1) spending more time exploring the obstacle; 2) running along the homing-vector path and then being blocked by the obstacle; 3) learning a subgoal at the other obstacle edge; and 4) targeting the obstacle edge while running away from the shelter. Furthermore, several minutes of exploration after obstacle removal was not sufficient to generate direct, homing-vector responses. The subgoal strategy is thus a hybrid between classical map-based ('locale') navigation, which results in an allocentric schema of space through exploration, and route-based ('taxon') navigation, in which learning occurs by repeating previous goal-directed trajectories (O'Keefe and Nadel, 1978).

Our results agree with ethology studies that have described movement patterns in the wild as highly dependent on the individual animal's prior paths (Benhamou, 1991; Meade et al., 2005; Thompson, 1982). Many strategies, such as odor trails, learning visual beacons, and learning fine-grained actions, could in principle underlie these previous results. Here we identify repeating subgoal targets as a fundamental, spontaneous learning process for mouse escape routes. In particular, our analysis of edge-vector movements supports the following working model: mice spontaneously execute visually guided edge-vector movements toward a salient wall edge during exploration or escape; if this movement brings the mouse toward a goal (here, the shelter), then

its target is memorized as a subgoal location that should be visited when escaping from behind the subgoal. We hypothesize that a one-shot learning rule works on the spatial targets of these instinctive actions, but further experiments should be done to directly test for causality.

Memorizing subgoals confers distinct survival advantages: it drives escape routes with the optimality of map-based planning and the rapidity of instinctive responses. However, this strategy is less flexible than sensory-guided or map-based mechanisms: mice continue to execute roundabout escapes in the absence of an obstacle, until re-establishing the homing-vector path. Responses to imminent predatory threats are known to favor quick reaction times at the expense of computational sophistication (Mobbs et al., 2020), so this strategy could in principle be specific to escape from imminent threat. Nonetheless, we found that it was also used in a less urgent food-seeking task, suggesting that it might be a general building block for quickly learning spatial locations that are important for survival. Indeed, unlike food-approach paths, exploratory movements across the arena were unaffected by experience with the obstacle. Thus, subgoals are not simply preferred locations but rather take part in specific goal-directed behaviors.

Subgoal learning bears some resemblance to hierarchical reinforcement learning (HRL), a technique in artificial intelligence for learning multi-step behaviors (Sutton et al., 1999). However, the learning process we have observed in mice does not fit cleanly into the dominant 'model-free vs. model-based' framework for reinforcement learning agents (cf. Spiers and Gilbert, 2015). Rather, it fuses action repetition with a model of space: mice discover a map of individual subgoals through targeted exploration and learning heuristics.

Finally, our results provide an entry point for studying the neural mechanisms of natural spatial learning in mice. Grounding explanatory frameworks for neuroscience in an understanding of the species' spontaneous behavioral repertoire promises to uncover new, foundational relationships between neural activity and behavior (Datta et al., 2019; Krakauer et al., 2017; Mobbs et al., 2018). It will be particularly informative to examine how systems such as the hippocampal formation and the midbrain interact during exploration and escape in order to produce instinctive navigational actions and subsequent learning of subgoals.

Experimental procedures

Animals

All experiments were performed under the UK Animals (Scientific Procedures) Act of 1986 (PPL 70/7652) following local ethical approval. We used singly housed (1 week), male, 8-12-week-old C57BL/6 mice during the light phase of the light cycle. All experiments were performed in naïve mice (no previous experimental experience), except for the 'experienced mice' condition of escaping in the dark and the wall length change experiments.

Behavioral arenas

The main arena was an elevated white acrylic circular platform 92 cm in diameter (Extended Data Fig. 8A). is the platform had a 51 cm x 1 cm hole in its center; through this hole, the obstacle (white acrylic, 50 cm long x 12.5 cm tall x .5 cm thick) could be raised (obstacle condition) or lowered (open field condition). For experiments in which the obstacle appears or disappears, this was done by digitally triggering a custom-made pneumatic tubing system (time to raise or lower the obstacle was ~100 ms). In the acute obstacle removal experiment, this was triggered simultaneously with the stimulus onset. In chronic obstacle removal experiments, this was triggered while the mouse was in the shelter. The hole obstacle consisted of a 50 cm long x 10 cm wide rectangular hole in the center of the arena (Extended Data Fig. 8B). The arena in the sideways-moving-obstacle experiment was an elevated white acrylic square platform 80 cm x 80 cm. The obstacle (white acrylic, 70 cm long x 12.5 cm high x .5 cm thick) was manually pulled to the "short" condition (obstructing 48 cm) or to the "long" condition (obstructing 66 cm) while the mouse was in the shelter. The shelter was a 10 cm cube of transparent red acrylic (opaque to the mouse). It included a mouse-hole-shaped entrance at the front and additional 2.5 cm tall square of red acrylic on top in order to prevent the mice from climbing on top. The arena was surrounded by a black, square, plastic surrounding. A projector screen was located above the arena. For the main experiments, light was projected onto the screen at 7.95 cd/m2. Note that the room was not totally sonically insulated and that neither the black surround nor the overhead illumination was circularly symmetric; these asymmetries could all provide orientation cues. The arena and shelter were cleaned with 70% ethanol after each session. The arena was illuminated with infrared light, and experiments were recorded at 30 frames-per-second with an infrared selective camera.

Escape behavior

Animals were given a 7-minute acclimation period during which they discovered the shelter. Stimuli were subsequently delivered when the mouse entered the threat zone (the back 20 cm, on the end opposite from the shelter) and was generally facing away from the shelter. Only stimuli delivered in this zone were analyzed. At least one minute was

allowed in between trials. Threat stimuli were loud (87 dB), unexpected crashing sounds played from a speaker located 1m above the center of the arena. Free sounds ("smashing" and "crackling fireplace") were downloaded from soundbible.com. They were then edited using Audacity software such that they were equally and continuously loud. Stimuli were alternated between the "smashing" sound and the "crackling" sound to prevent stimulus habituation. In some sessions (4 with and 4 without an obstacle), we used an ultrasonic sweep stimulus (17-20 kHz, 3 sec). No difference in response between the stimuli was observed and therefore the data from these sessions were pooled. For each trial, the stimulus was triggered repeatedly until the mouse reached the shelter, for a maximum 9 seconds. Since escapes take longer with the hole obstacle and in the dark, stimuli in these conditions were played for up to 12 seconds. Stimulus responses were considered as escapes if the mouse reached the shelter within 12 seconds in the light or 18 seconds for the hole obstacle and for the wall obstacle in the dark. Mice varied in how many trials they performed in each experiment. We thus limited analysis to the first 3 escapes in each condition (more than 50% of mice completed at least 3 escapes in all experiments). In the arena with the wall obstacle, the probability of threat-evoked escape to shelter was 93%, while the probability with a no-stimulus control was 12%.

Food-approach behavior

Mice were food restricted to 85% of their baseline weight. They were then trained in 5, ~1-hour sessions to approach and lick a metal spout in response to a 9-second 10-kHz tone. Training was done in a 60cm x 15cm rectangular arena. Licking the spout while the tone was on resulted in a 7-µL drop of condensed milk (diluted 1:1 with water) being delivered. For the lick raster plots, licks were plotted at 5 licks/sec when the sensor was tonically triggered by licking. These mice had two sessions. The first session was in the escape arena with no obstacle, the shelter in its usual location, and a lick spout on the opposite end of the arena. They received "practice trials" of tone and milk, mostly when they were near the lick spout. After 20 minutes, trials were initiated when the mouse was on the opposite side from the lick spout, and these data were used for analysis. In the second session, the obstacle was initially present, and then after 20 minutes it was removed, while the mouse was in the shelter. Mice performed more trials than with the escape behavior, so here we examined trajectories from the first 9 successful trials (greater than 50% of mice completed at least 9 trials). For testing correlations, to match the analysis performed on escapes, we limited analysis to the first 3 trials.

Video tracking and visualization

Videos were acquired at 30 frames per second using an overhead camera with a near-infrared-selective filter. Videos were then fisheye-distortion corrected, aligned onto a common coordinate framework, and visualized with custom Python code using the OpenCV library (github.com/BrancoLab/Common-Coordinate-Behaviour; Supplementary Video 1). We used DeepLabCut (github.com/AlexEMG/ DeepLabCut) to track the mouse from the video, after labelling 1500

frames with 13 body parts. Post-processing includes removing low-confidence tracking, using a median filter to discard outliers, and applying an affine transformation to the tracked coordinates to match the common coordinate framework.

Analysis

All analysis was done using custom software written in Python (github.com/BrancoLab/escape-analysis). Escape initiation time was computed as the time from the stimulus onset to the first moment that the center of the mouse moves toward the shelter at ≥62.5 cm s-1 averaged over 3 frames. The initial escape direction was computed by taking the vector from the mouse's position at the reaction time to its position when it is 10 cm in front of the obstacle. We computed a direction score where a vector aimed directly at the shelter received a value of 0; one aimed at either obstacle edge received a value of 1.0; a vector halfway between these would score 0.5; and a vector that points beyond the edge would receive a value greater than 1.0. The formula is: direction= $abs((offset_{HV}-offset_{EV}+offset_{HV-EV})/(2 * offset_{HV-EV}))$, where $offset_{HV}$ is the distance from the mouse to where the mouse would be if it took the homing vector, $offset_{EV}$ is the distance from the mouse to where the mouse would be if it took the obstacle-edge vector, and $offset_{HV-EV}$ is the distance from the homing-vector path to the obstacle-edge-vector path. Only the obstacle edge closest to the escape path was considered. Initial food-approach trajectories and spontaneous exploration trajectories were analyzed in the same manner. The threshold for classifying a trajectory as an obstacle edge vector (as opposed to a homing vector) was the 95th percentile of escapes in the open-field condition (0.68). For examining the effect of experience on spontaneous exploration, we used a threshold of 0.5 to distinguish center-directed and edge-directed movements.

Previous homings used to predict escape trajectories included both stimulus-evoked escapes and spontaneous homings. These were extracted by taking movements that start in the side of the arena opposite the shelter and behind the wall by ≥10 cm, and then move toward the shelter or obstacle edges without any pauses longer than 1 s or moments in which the mouse faces away from the shelter direction. To obtain the mean direction of previous movements, we averaged the 10 most recent homings and then computed the direction score using the formula above. To predict escape direction, we needed to prevent previous obstacle-edge vectors on the right from influencing our prediction of whether the mouse will execute an obstacle-edge vector on the left and vice versa. To achieve this, we did not consider previous movements targeting the left 20% of the obstacle (10 cm) when making a prediction for the trajectory of escapes that are closer to the right edge.

Spontaneous exploratory traversals were considered exploratory paths that start at either end of the arena (within 20 cm of either end) and then cross the center. Traversals that go along the boundary of the arena (i.e. greater than 10 cm past the obstacle edge) or take longer than 2 seconds were excluded from analysis.

For permutation tests, the test statistic is the group mean difference (e.g. in escape direction or path efficiency). The condition of each mouse (e.g. open field vs obstacle) is randomly shuffled 10,000 times to generate a null

distribution. For comparisons on quantities other than direction, outliers (> 2 IQR from median) were removed. Tests for differences in efficiency, reaction time, and initial escape conditions were two-tailed; tests for increases in edge-vector responses compared to an open-field control were one-tailed. The ANOVA was performed using the linear mixed effects model package in R, after removing outliers (z-score > 0.975).

For the linear prediction models, we used ridge regression with regularizer alpha = 0.1. Predictive power on escape direction using exploration or turn angles was not dependent on alpha in the range [0, 10], while predictive power using prior movements was constant in the range [0, 0.5] and gradually decreased above alpha = 0.5. 1,000 shuffles of a 2-fold cross validation were performed, and the probability density of r^2 prediction scores on the test data was plotted. To predict using prior escape directions, we used the direction of the mean target of prior homing movements, as described above. To subtract the contribution of time spent exploring to prior homing movement targets and vice versa, we performed a linear regression predicting each variable from the other. We then used the residual values from this prediction as input to the main linear prediction model. To weight prior homing movements by their initial position or angle, we used a weighted mean target of prior homing movements in the prediction model. The weight of each prior movement varied between 0 (initial positions ≥50 cm apart or body angle off by 180º) and 1 (same initial position or body angle); to weight based on how long it takes to end up in the shelter after for executing a homing movement from the threat zone, weights varied between 0.1 (≥12 seconds) and 1 (0 seconds). In all weightings, weights were normalized to sum to 1.

To predict using turn angles, escapes and spontaneous movements were first decomposed into individual turn-and-run movements in which the mouse continuously turns and/or runs in a consistent direction. Only movements in which the mouse moved at least 10 cm closer to the shelter or obstacle edges were used. We used the angle between the mouse's body at the starting point of the movement and its body angle after it had moved 10 cm away from that point. This value is negative for left turns and positive for right turns. For predicting direction from turn angles, we trigonometrically extrapolated where the mouse would end up if it followed the predicted turn angle, and computed initial escape direction as above. The data that support the findings of this study are available from the corresponding authors upon request.

Acknowledgments

This work was funded by a Wellcome Senior Research Fellowship (214352/Z/18/Z) and by the Sainsbury Wellcome Centre Core Grant from the Gatsby Charitable Foundation and Wellcome (090843/F/09/Z) (T.B.), and the SWC PhD Programme (P.S. and S.O.). We thank members of the Branco lab and T. Mrsic-Flogel for discussions; J. Rapela for advice on statistical analysis; T. Mrsic-Flogel, T. Behrens, C. Barry, M. Stephenson-Jones, Y. Isogai, C. Clopath, Y.

Lin Tan, F. Claudi and R. Vale for comments on the manuscript; the SWC Neurobiological Research Facility and FabLabs for technical support; K. Betsios for programming the data acquisition software.

References

Alyan, S., Jander, R., 1994. Short-range homing in the house mouse, Mus musculus: stages in the learning of directions. Animal Behaviour 48, 285–298. https://doi.org/10.1006/anbe.1994.1242

Benhamou, S., 1991. An analysis of movements of the wood mouse Apodemus sylvaticus in its home range. Behavioural Processes 22, 235–250. https://doi.org/10.1016/0376-6357(91)90097-J

Collett, T.S., 1982. Do toads plan routes? A study of the detour behaviour ofBufo viridis. J. Comp. Physiol. 146, 261–271. https://doi.org/10.1007/BF00610246

Cooper Jr., W.E., Blumstein, D.T. (Eds.), 2015. Escaping From Predators: An Integrative View of Escape Decisions. Cambridge University Press, Cambridge. https://doi.org/10.1017/CBO9781107447189

Datta, S.R., Anderson, D.J., Branson, K., Perona, P., Leifer, A., 2019. Computational Neuroethology: A Call to Action. Neuron 104, 11–24. https://doi.org/10.1016/j.neuron.2019.09.038

Drickamer, Lee.C., Stuart, J., 1984. Peromyscus: Snow Tracking and Possible Cues Used for Navigation. American Midland Naturalist 111, 202. https://doi.org/10.2307/2425561

Edvardsen, V., Bicanski, A., Burgess, N., 2020. Navigating with grid and place cells in cluttered environments. Hippocampus 30, 220–232. https://doi.org/10.1002/hipo.23147

Ellard, C.G., Eller, M.C., 2009. Spatial cognition in the gerbil: computing optimal escape routes from visual threats. Animal Cognition 12, 333–345. https://doi.org/10.1007/s10071-008-0193-9

Etienne, A.S., Jeffery, K.J., 2004. Path integration in mammals. Hippocampus 14, 180–192. https://doi.org/10.1002/hipo.10173

Etienne, A.S., Teroni, E., Maurer, R., Portenier, V., Saucy, F., 1985. Short-distance homing in a small mammal: the role of exteroceptive cues and path integration. Experientia 41, 122–125. https://doi.org/10.1007/BF02005909

Freas, C.A., Schultheiss, P., 2018. How to Navigate in Different Environments and Situations: Lessons From Ants. Front. Psychol. 9. https://doi.org/10.3389/fpsyg.2018.00841

Hamilton, D.A., Rosenfelt, C.S., Whishaw, I.Q., 2004. Sequential control of navigation by locale and taxon cues in the Morris water task. Behavioural Brain Research 154, 385–397. https://doi.org/10.1016/j.bbr.2004.03.005

Harrison, F.E., Reiserer, R.S., Tomarken, A.J., McDonald, M.P., 2006. Spatial and nonspatial escape strategies in the Barnes maze. Learning & Memory 13, 809–819. https://doi.org/10.1101/lm.334306

Janson, C., 2014. Death of the (traveling) salesman: Primates do not show clear evidence of multi-step route planning: Lack of Complex Route Planning. Am. J. Primatol. 76, 410–420. https://doi.org/10.1002/ajp.22186

Krakauer, J.W., Ghazanfar, A.A., Gomez-Marin, A., MacIver, M.A., Poeppel, D., 2017. Neuroscience Needs Behavior: Correcting a Reductionist Bias. Neuron 93, 480–490. https://doi.org/10.1016/j.neuron.2016.12.041

Layne, J.E., 2003. Mechanisms of homing in the fiddler crab Uca rapax 1. Spatial and temporal characteristics of a system of small-scale navigation. Journal of Experimental Biology 206, 4413–4423. https://doi.org/10.1242/jeb.00660

Lima, S.L., Dill, L.M., 1990. Behavioral decisions made under the risk of predation: a review and prospectus. Can. J. Zool. 68, 619–640. https://doi.org/10.1139/z90-092

Liu, A., Papale, A.E., Hengenius, J., Patel, K., Ermentrout, B., Urban, N.N., 2020. Mouse Navigation Strategies for Odor Source Localization. Front. Neurosci. 14, 218. https://doi.org/10.3389/fnins.2020.00218

McMillan, B.R., Kaufman, D.W., 1995. Travel path characteristics for free-living white-footed mice (*Peromyscus leucopus*). Can. J. Zool. 73, 1474–1478. https://doi.org/10.1139/z95-174

Meade, J., Biro, D., Guilford, T., 2005. Homing pigeons develop local route stereotypy. Proceedings of the Royal Society B: Biological Sciences 272, 17–23. https://doi.org/10.1098/rspb.2004.2873

Mobbs, D., Headley, D.B., Ding, W., Dayan, P., 2020. Space, Time, and Fear: Survival Computations along Defensive Circuits. Trends in Cognitive Sciences 24, 228–241. https://doi.org/10.1016/j.tics.2019.12.016

Mobbs, D., Trimmer, P.C., Blumstein, D.T., Dayan, P., 2018. Foraging for foundations in decision neuroscience: insights from ethology. Nat Rev Neurosci 19, 419–427. https://doi.org/10.1038/s41583-018-0010-7

Morris, R.G.M., 1981. Spatial localization does not require the presence of local cues. Learning and Motivation 12, 239–260. https://doi.org/10.1016/0023-9690(81)90020-5

O'Keefe, J., Nadel, L., 1978. The hippocampus as a cognitive map. Clarendon Press ; Oxford University Press, Oxford : New York.

Restle, F., 1957. Discrimination of cues in mazes: A resolution of the "place-vs.-response" question. Psychological Review 64, 217–228. https://doi.org/10.1037/h0040678

Simon, P., Dupuis, R., Costentin, J., 1994. Thigmotaxis as an index of anxiety in mice. Influence of dopaminergic transmissions. Behavioural Brain Research 61, 59–64. https://doi.org/10.1016/0166-4328(94)90008-6

Spiers, H.J., Gilbert, S.J., 2015. Solving the detour problem in navigation: a model of prefrontal and hippocampal interactions. Front. Hum. Neurosci. 9. https://doi.org/10.3389/fnhum.2015.00125

Stachenfeld, K.L., Botvinick, M.M., Gershman, S.J., 2017. The hippocampus as a predictive map. Nature Neuroscience 20, 1643–1653. https://doi.org/10.1038/nn.4650

Sutton, R.S., Precup, D., Singh, S., 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. Artificial Intelligence 112, 181–211. https://doi.org/10.1016/S0004-3702(99)00052-1

Teichroeb, J.A., Smeltzer, E.A., 2018. Vervet monkey (Chlorocebus pygerythrus) behavior in a multi-destination route: Evidence for planning ahead when heuristics fail. PLoS ONE 13, e0198076. https://doi.org/10.1371/journal.pone.0198076

Thompson, S.D., 1982. Spatial Utilization and Foraging Behavior of the Desert Woodrat, Neotoma lepida lepida. Journal of Mammalogy 63, 570–581. https://doi.org/10.2307/1380261

Tolman, E.C., 1948. Cognitive maps in rats and men. Psychological Review 55, 189–208. https://doi.org/10.1037/h0061626

Vale, R., Evans, D.A., Branco, T., 2017. Rapid Spatial Learning Controls Instinctive Defensive Behavior in Mice. Current Biology 27, 1342–1349. https://doi.org/10.1016/j.cub.2017.03.031

Wallace, D.G., Gorny, B., Whishaw, I.Q., 2002. Rats can track odors, other rats, and themselves: implications for the study of spatial behavior. Behavioural Brain Research 131, 185–192. https://doi.org/10.1016/S0166-4328(01)00384-9

Yilmaz, M., Meister, M., 2013. Rapid Innate Defensive Responses of Mice to Looming Visual Stimuli. Current Biology 23, 2011–2015. https://doi.org/10.1016/j.cub.2013.08.015