A robot-rodent interaction arena with adjustable spatial complexity for ethologically relevant behavioral studies

Alexander T. Lai[a,*], German Espinosa[b,*], Gabrielle E. Wink[c,*], Christopher F. Angeloni[d,*], Daniel A. Dombeck[d,**], Malcolm A. MacIver[a,b,c,d,**]

[a]*Northwestern University, Department of Biomedical Engineering, Technological Institute E311, 2145 Sheridan Road, Evanston, 60208, IL, USA*
[b]*Northwestern University, Department of Computer Science, Seeley Mudd 3219, 2233 Tech Drive, Evanston, 60208, IL, USA*
[c]*Northwestern University, Department of Mechanical Engineering, Technological Institute B224, 2145 Sheridan Road, Evanston, 60208, IL, USA*
[d]*Northwestern University, Department of Neurobiology, Hogan 2-160, 2205 Tech Drive, Evanston, 60208, IL, USA*

Abstract

Outside of the laboratory, animals behave in spaces where they can transition between open areas and coverage as they interact with others. Replicating these conditions in the laboratory can be difficult to control and record. This has led to a dominance of relatively simple, static behavioral paradigms that reduce the ethological relevance of behaviors and may alter the engagement of cognitive processes such as planning and decision making. Therefore, we developed a method for controllable, repeatable interactions with others in a reconfigurable space. Mice navigate a large honeycomb lattice of adjustable obstacles as they interact with an autonomous robot coupled to their actions. We illustrate the system using the robot as a pseudo-predator, delivering airpuffs to the mice. The combination of obstacles and mobile threat elicits a diverse set of behaviors—such as increased path diversity, peeking, and baiting—providing the foundations to explore ethologically relevant behaviors in the laboratory.

Introduction

The rich emergent behaviors that neuroscience seeks to understand occur in natural environments in which there is variability in cover, for example, from open areas to more cluttered spaces [1, 2], and in which competition or cooperation with other animals occurs. Such environmental variability and interactivity is absent from most laboratory paradigms for rodents, even though the neural circuits driving behavior likely evolved for survival in these conditions. Here we describe an experimental system which attempts to encourage more ethological behaviors by combining two rarely combined features: a spatially complex arena and an interactive agent.

*Equal contribution author
**Corresponding author
Lead contact: Malcolm A. MacIver

With some recent exceptions [3, 4, 5, 6], traditional laboratory arenas are static and non-interactive, with appetitive or aversive inanimate stimuli provided at fixed times or locations, reducing task complexity. These studies have revealed a rich array of cognitive representations of the latent variables describing behavior in such environments, such as place, heading direction and grid cells, and neurons storing choice or value information [7, 8, 9, 10, 11]. However, it is possible that cognitive representations such as these are engaged differently in more ethologically relevant conditions. For example, recent research has begun to address the question of how the brain encodes the location and behavioral tendencies of others, but these studies were largely performed in conditions where the other was not task relevant [12, 13]. Few if any place cells of others were identified, possibly because of the lack of task relevance of the other. More recent experiments increasing the task relevance of the other [14, 15] have shown that self-place cell firing can be modulated by the location of conspecifics in the environment. Since these circuits likely evolved to encode ethologically relevant interactions, a greater understanding of their function is likely to emerge as experiments approach more natural conditions, highlighting the need for new, more ethologically relevant laboratory paradigms for studying interactions with others.

Of course, a major advantage of the sparse and simple spatial layout of traditional laboratory arenas, such as open field, linear tracks and T-mazes [16, 17, 18], is the ability to perform highly repeatable, controlled experiments which maximize statistical power. Intuitively, the spatial complexity of these spaces differs considerably from that of natural environments. We provide quantitative evidence for this difference below. It is possible that the simplicity of traditional experimental paradigms alters cognitive processing in animals behaving in these spaces, or make particular processes difficult to study. For example, the neural substrates of planning have not been clearly established. Many studies have investigated this question [19, 20, 21], though largely in traditional simple mazes. One of the most likely substrates are the "preplay" events in the hippocampus during sharp wave ripples, which lead to rapid sequential activation of remote place cells. But large debates persist about whether these neural signals represent recall of past trajectories or are in fact thoughts about the future [22]. This problem is exacerbated by what could be called the 'Groundhog Day effect' of highly simplified spatial layouts: If the space an animal has experienced in the past is unchanged from the one that it will experience in the future, then it is difficult or impossible to disentangle memory from foresight. Notably, using a task which increased trial-to-trial path diversity provided some of the best evidence for planning, with prospective replay events often seemingly predicting future navigation paths [19].

A task logic analog of the Groundhog Day effect is that current task designs result in the test subjects quickly learning the task contingencies, leading to habituation and a reduction of behavioral indicators of planning such as VTE [23]. Thus, the statistical consistency of an animal's path through a typical laboratory test environment, or the repetitious nature of the task itself, appears to be important variables in the study of cognitive processes such as planning. Yet, it is rarely varied systematically in experiments, particularly to the level found in more natural contexts. Given these con-

siderations and broader calls for new laboratory paradigms to probe animal behavior in complex, ethologically relevant scenarios [24, 25, 26], there is an opportunity to bring some of the complexity of natural scenarios into the laboratory without compromising on the control of the experiment's variables and statistical power.

To address these issues, we designed a system which provides the control and repeatability of previous paradigms, yet facilitates more naturalistic behavioral tasks through two key innovations: rapidly reconfigurable obstacles and a mobile robotic agent (Fig. 1). The physical basis of the system is an arena with removable obstacles in a honeycomb lattice: this allows the experimenter to vary spatial complexity, enabling configuration in naturalistic partially cluttered arrangements, and facilitates rapid switching between spatial layouts. Multiple high speed cameras ensure reliable tracking of mice throughout the space despite these obstacles. Controllable interaction with an "other" is provided by a mobile, wireless robot that is coupled to the behavior of the mouse with negligible latency. Finally, automation allows multiple hours of operation without human intervention beyond animal subject and robot battery replacement.

Here, we provide details on the design and implementation of this system, termed cellworld for brevity, and discuss results from one particular implementation which emulates naturalistic predator-prey encounters by pitting the mouse against an airpuff-equipped predator-like variant of the robot. Several other possible configurations, such as using the robot as prey or for phonotactic localization, are described in Supplementary Table 1.

With the robot-predator configuration, we found evidence of a rich array of behaviors, spanning from highly variable trajectories (occurring at a rate of ≈1 trial/min or slower) to trajectories that are used repeatedly with little variation (occurring at a rate of ≈2 trial/min or faster). High variability trajectories included peeking at and seemingly luring or baiting the robot predator away from the location the mouse needed to reach for its reward, not unlike the broken-wing display found in birds [27]. Peeking in rodents appears to emerge in the context of more complex naturalistic conditions [28, 29, 30, 31], and we are not aware of prior observations of baiting in rodents. These behaviors and the path diversity that we observed across trials may be specifically useful for future studies into the cognitive representation of others or mechanisms of planning; and cellworld may be generally useful for enriching task designs for research into decision making, navigation, learning, memory, fear, and anxiety, among other domains (Supplementary Table 1).

Results

Creating naturally inspired spaces with a reconfigurable arena

The hexagonal arena, shown in Figure 2, is 2.34 m at its widest length (2.56 m^2), and is comprised of 331 hexagonal cells, with a center-to-center distance of ≈11 cm, slightly more than one standard adult mouse body length. Each cell has a pair of magnets for securing of obstacles above a thin vinyl membrane (for removal of odor cues between subjects) that lies on top of the magnetized floor. This design allowed us to rapidly reconfigure the arrangement of obstacles for each experimental session (Fig. 2b-d).

With the goal of studying behavior in a more ethological context, we used a genera-

tive algorithm to create arenas that more closely resembled the spatial statistics of natural landscapes. To accomplish this, we used a single parameter, entropy[1], to create random arrangements of obstacles (Fig. 2e). In the simplest terms, entropy describes the degree of clutter in a space, such that a space with very few obstacles has low entropy and a space that is half-filled with obstacles has maximum entropy. Next we measured how these more naturalistic spaces compared to classical laboratory setups for studying rodent behavior. To accomplish this, we recreated classical mazes from prior studies, including linear tracks, T-mazes and radial arm mazes (Fig. 2d).

We hypothesized that the complexity of the experimental space might be useful for natural behaviors, and therefore considered the visual connectedness of various arena layouts. To do so, we computed the network degree complexity (hereafter spatial complexity[1]) of generated arenas, our versions of classical mazes, and other spaces such as natural landscapes. Spatial complexity summarizes the visual connectedness of a space: high complexity arenas contain a mix of short and long sightlines, while low complexity arenas contain primarily short or primarily long sightlines. Intuitively, this measure relates to the behavioral utility of a space: high complexity spaces provide a mix of hiding spots and long sightlines to gather information, features which may be useful for evading a predator or planning.

We generated 500 random arena configurations at 14 different entropy levels (*Methods*) and then computed spatial complexity for each of these arenas. We found that the spatial complexity of the arenas peaked at mid-range levels of entropy (0.4–0.5) with a complexity of 0.80 $\pm$ 0.02 (Fig. 2e). This value is similar to the most prevalent spatial complexity value found by repeatedly sampling satellite images of a savanna landscape (0.80, see *Methods*; Fig. 2e, right panel). Spatial complexity analyses of other savanna samples and key terrestrial habitats have similar results (Mugan & MacIver[1], Supplementary Fig. 11 and Supplementary Table 1). In comparison, some complexities of our renderings of classical mazes were found to be much lower than these natural landscapes, ranging between 0.00–0.17. These results suggest that by controlling the entropy level of randomly generated obstacles, we can control the complexity of the cellworld arena. Furthermore, arenas generated with mid-level entropy are more similar to natural landscapes than to classical maze designs. Based on these results, we hypothesized that a subset of the generated arenas are ideal for planning and evasion, and therefore focused our later behavioral experiments on the two extremes of spatial complexity: an open arena (entropy: 0.0, spatial complexity: 0.0) and an obstacle configuration with mid-level entropy (entropy: 0.5, spatial complexity: 0.74; Fig. 2c, middle panel). However, these spatially complex environments contain a large number of occlusions, requiring a multi-view tracking system for consistent behavioral monitoring, which we describe next.

A multi-view camera system for continuous tracking in occluded spaces

We designed the camera system in cellworld to meet two experimental goals: 1) to consistently observe the mouse's position in spatially complex arenas, and 2) to control the behavior of a mechanical agent with negligible latency after automatically detected changes in mouse position and orientation (which we describe in the next sec-

tion). To meet these goals, the system uses four high frame rate and low-latency infrared cameras. The cameras are suspended 200 cm above the arena floor and each cover a specific quadrant (Fig 3a), capturing 2040 x 2040 pixel images at a rate of 120 frames per second (fps). This layout is designed to minimize blind spots created by obstacles within the arena—a crucial aspect as important behaviors could occur near these obstacles (Fig 3b). Additionally, the high frame rate and low latency of these cameras enabled real-time monitoring of animal movements, allowing us to couple the behavior of an autonomous robot to that of the mouse.

To perform mouse tracking, we acquired perspective-corrected, stitched images from the four cameras (*Methods*), then removed all static elements using background subtraction. The remaining features (mouse, robot) were identified using color-connected components. Robot tracking was simplified through three LEDs on the top of the robot (Fig 3f). This enabled us to perform real-time monitoring of robot and mouse movements with an average latency of 3.2 ms and facilitated swift response to changes in animal behavior. For the current study, the frame rate and throughput of the system was capped at 90 Hz as that was found to be sufficient for updating the robot's heading when moving quickly through obstacle fields, but cellworld's tracking system can process a maximum of 206 fps (Supplementary Fig. 2b). In summary, this tracking system allowed continuous behavioral monitoring of a mouse in a densely occluded, ethologically-inspired space, while also facilitating low latency control of a robotic agent.

An autonomous mobile agent coupled to animal behavior

A crucial aspect of natural behavior in many species is interaction with others, but these behaviors can be difficult to control in the lab. To that end, we engineered a fully autonomous robot (Fig. 3c-e) whose behavior is tied to that of the mouse with no more than 11 ms of latency. The robot itself has no vision system, but we synthesized an omnidirectional visual sensory volume based on images from the camera system, the location of the robot, and the location of the obstacles. We then controlled the robot in closed-loop to pursue the mouse when it was in "view" and to otherwise search unseen locations when the mouse was out of view (Fig. 4a). (Note that this means freezing responses on the part of the rodent have no effect on the robot's ability to perceive; it would be simple to modify this such that the robot only "sees" the mouse upon movement.) Next, we took advantage of this low-latency coupling between the robot and the mouse's behavior to simulate predator-prey interactions in the lab.

To do so, we outfitted the robot with an airpuff module, which consisted of a CO_2 tank and valve actuated via a motor to release a sequence of two brief, powerful blasts of air when the mouse came within 32 cm of the robot (Fig. 4a, Supplementary Movies 1, 2). We term this aversive airpuff sequence an "attack" event, but note that due to the modular design of the robot, other stimulus modes (such as appetitive rewards, visual, or auditory stimuli) may be used.

To test whether the ability to attack made the robot more behaviorally relevant to the mouse, we performed a pilot study where mice first interacted with a stationary or pursuing robot with the airpuff disabled, then enabled the airpuff for the following ex-

perimental sessions. Consistent with prior airpuff results [32, 33], we found that the mice significantly increased the distance between their location and the robot after entering the attack threshold (n=2 mice; airpuff enabled: 93.0 ± 36.1 cm, airpuff disabled median ± IQR distance: 25.5 ± 10.7 cm, $p = 8.16 \times 10^{-78}$; Supplementary Fig. 3, Supplementary Movies 3, 4). From this we concluded that the airpuff-equipped robot was behaviorally relevant to the mouse, and induced escape or avoidance behaviors, which allowed us to leverage cellworld to create a task inspired by predator-prey dynamics.

A predator-prey inspired behavioral task disrupts stereotyped navigation

With the capability of creating a spatially complex arena patrolled by an aversive robot, we devised a behavioral task modeled on predator-prey interactions. In this task, mice start on one side of the arena, and must navigate to the other side of the arena while evading a pursuing robot to reach a water reward (the robot evade task, or BotEvade hereafter for brevity, Fig. 4).

In order to facilitate multiple mouse traversals within a single 30 minute experimental session, we engineered several additional components for the cellworld system: chambers containing water lick ports and automated doors at the start and end of the arena, and an external return chute that connects the chambers to allow the mouse to return to the start and reset the system for another trial. Both the doors and the water feeders were controlled and monitored by software that coordinates events between the robot, rodent, and components of the arena, termed the "experiment controller," which used lick events detected by the water feeders to determine when to start and end trials. Mice were guided through the task by specific sequences of door events (Fig 4**b, c**), allowing them to initiate and complete many trials per experiment (the maximum completion rate for n=8 mice was 83 ± 10 trials per 30 minute session). With these systems in place, we developed a training protocol to encourage repeated interaction between the mouse and the robot within a spatially complex arrangement of obstacles.

To do so, we trained mice in the BotEvade task with the following steps. First, 8 mice were acclimated to the reward zones and return chute using a gutter-like corridor directly linking the entrance and exit of the arena (CT: corridor training). Next, the corridor was removed, and obstacles were placed in a mid-entropy arena (the Random 0.5 arena of Fig. 2c) to allow the mice to learn the spatial layout (T: arena training phase). Once the mouse behavior stabilized, the robot was introduced to the environment (R: robot phase). Then, once behavior in the presence of the robot stabilized, we removed the robot from the arena to measure extinction of the behavioral response to the autonomous predator (PR: post-robot phase; Fig. 5b). Mice learned the task rapidly, taking 4.0 ± 2.1 days to plateau during the T phase (Fig. 5c). We also found that the airpuff equipped robot was an effective aversive stimulus, eliciting fleeing behaviors in 74.7% airpuff events when compared to shuffled data (n = 178 attack events; Fig. 5d, Supplementary Movie 2). Thus, we found that our training protocol encouraged mice to repeatedly traverse a spatially complex environment, creating numerous interactions with the aversive robot over the course of the experiments.

We predicted that the combination of a spatially complex layout and predatory agent would elicit a richer set of behaviors compared to a simple spatial layout with-

out a predatory agent. To directly assess the effect of these two variables, we compared the cohort of mice in the mid-entropy arena to an additional cohort of mice which was trained in an open field (in these experiments, the PR phase was omitted). During training (T phase), mice had highly variable trajectories with (n=8) or without obstacles (n=2) as they explored the environment and learned the task (Fig. 6a, left column, Supplementary Movie 5). When the robot was introduced (R phase), different behavioral patterns emerged in the two environments: in the occluded arena, routes became more variable and slower, while in the open arena, mice reverted to thigmotaxis— running along either the north or south wall at high speeds (Fig. 6a, middle column, Supplementary Movie 6). Interestingly, when the robot was removed from the occluded arena (PR phase), mice largely reverted to two thigmotactic routes along the north and south walls of the arena (Fig. 6a, right column, Supplementary Movie 7).

The highly variable routes in the presence of the robot in the spatially complex arena suggested that mice engaged in more sophisticated evasion strategies in complicated environments, therefore, we focused our subsequent analyses on these experiments. We found that mice completed significantly fewer trials per 30 minute session in the R phase (27.0 ± 14.6 trials) than in the PR phase (57.6 ± 20.5 trials, $p = 1.44 \times 10^{-4}$; Fig. 6b), taking significantly longer routes to reach the goal during the R phase (420 ± 33 cm; 1.8 times the shortest path length of 234 cm) compared to the PR phase (340 ± 19 cm; 1.4 times the shortest length, $p = 0.002$; Fig. 6c).

We suspected this increase in route length occurred because, 1) the mice chose new routes after being exposed to the robot and, 2) when encountering the robot along a preferred route, mice changed course to evade it. To test these two hypotheses we used QuickBundles[34] to cluster the trajectories from each mouse in each experimental phase. To quantify the tendency to choose new routes, we counted the number of clusters found in each phase, and to quantify the tendency to deviate from a route, we calculated the average distance of each trajectory from the center of the nearest cluster (Fig. 6d). We found that there were significantly more clusters in the R phase (4.0 ± 0.5 clusters) than in the PR phase (1.0 ± 1.0 clusters, $p = 2.61 \times 10^{-4}$; Fig. 6e), and that trajectories tended to be further away from the nearest cluster in the R phase (17.6 ± 3.1 cm) compared to the PR phase (8.1 ± 3.3 cm, $p = 6.63 \times 10^{-5}$; Fig. 6f). Taken together, these results suggest that mice chose novel routes and deviated from preferred routes in order to evade the robotic threat.

Finally, we observed that mouse traversals were significantly slower in the R phase (68.2 ± 26.8 cm/s) compared to the PR phase (114.8 ± 26.2 cm/s, $p = 0.002$; Fig. 6g). This could reflect deceleration during rerouting, suggested by previous results (Fig. 6e-f), or it could reflect slow downs and stops. To test this, we quantified periods of time when mice paused during the experiments (*Methods*). We observed that mice paused more frequently near the entrance during the R phase (2.8 ± 1.8 pauses per trial) compared to the PR phase (1.2 ± 0.7 pauses per trial, $p_{\text{adj}} = 0.004$; Fig. 6h). Upon entering the arena, mice paused for longer durations in the R phase (1.5 ± 0.3 s) compared to the PR phase (1.1 ± 0.3 s; $p_{\text{adj}} = 0.017$; Fig. 6i). We also examined the frequency of pauses longer than 2 s in duration (Fig. 6i, *inset*), and found that longer pauses were more prevalent dur-

ing the R phase (0.3 ± 0.1 pauses per trial) compared to the PR phase (0.1 ± 0.1 pauses per trial, $p = 0.014$; Fig. 6j). Together, these results indicate that mouse behavior is significantly changed in the presence of the robot. Mice paused more at the arena entrance, suggesting they are more hesitant to enter, and they paused more frequently and for longer durations once inside the arena, possibly in order to hide or gather information about the robot location.

In summary, we used cellworld to assess mouse behavior in a spatially complex arena while interacting with an aversive "other" agent in the form of an airpuffing robot. We found that this combination of experimental features resulted in the disruption of habitual behavioral strategies, such as thigmotaxis and route stereotypy, and also resulted in increased pauses within the arena. Previous work has shown that such features may indicate planning [23], and we suspected that mice were using sequences of pauses to evade the robot. To assess this, we more closely examine examples of complex behaviors that we observed during the BotEvade task in the following section.

Presence of a robot in a spatially complex environment elicits complex behaviors

Above we established that mice took longer, more diverse paths at lower speeds when the robot was present and paused more often when engaging with the robot in the arena (Fig. 5). We hypothesized that these changes might reflect deliberation, such as monitoring the robot's movements to predict its future location, or planning new routes to evade the robot and reach the exit. We found several examples of behaviors consistent with this hypothesis. For instance, we found that mice engaged in apparent "baiting" behaviors, in which the mouse made visual contact with the robot, returned to a safe location (typically near the entrance), and then waited for the robot to approach. Once the robot approached the mouse, the mouse escaped along an open path opposite the robot's location (typically along the north wall) towards the exit (Fig 7a, Supplementary Movie 8), effectively leveraging their higher speed over that of the robot (speed of robot: 24.1 ± 1.2 cm/s, mouse: 68 ± 26.8 cm/s).

We also observed many instances of what appeared to be "peeking" behaviors. In one example (supplementary Movie 9), the mouse ran along a familiar path, then encountered the robot blocking the exit door (Fig 7b, left panel). After retreating to a safe location, the mouse then paused between two obstacles and centered the robot within its binocular zone while concealing its body behind a nearby obstacle (Fig. 7c). After seeming to confirm the robot's new location, the mouse then rerouted to a safe path avoiding the area near the robot and reached the arena exit and water reward (Fig 7b, right panel).

Instances of both baiting and peeking behaviors were found in all 8 out of 8 mice. While baiting is specific to the presence of the robot, we observed peeking in both the presence and absence of the robot (R and PR phases). In support of this, we manually identified 15 trajectories across R and PR phases where peeking events occurred and 10 trajectories from the R phase where baiting occurred. Movies of these trajectories can be accessed through the zenodo link found in the Key Resource Table under "Additional movies". Though these are only a subset of the many occurrences of these behaviors that we observed, they are a representative sample. Trajectories with

peeking events exhibited higher levels of deviation from typical trajectories, slower moving speed in the PR phase, and higher episode trajectory lengths and cluster distances in both the R and PR phases (Supplementary Fig. 4). Similarly, trajectories with baiting events had higher episode trajectory lengths and distances to the nearest cluster (Supplementary Fig. 4). All trajectories for each of the 10 baiting and 15 peeking trajectories are shown in Supplementary Figure 5 and 6. In addition, we have taken some initial steps to quantify these behaviors, plotting distances and visual contact between the prey and robot during the R phase "peeking" and "baiting" trajectories (Supplementary Fig. 5 and 6).

Taken together, these results show that the combination of a spatially complex arena and aversive robotic agent resulted in a rich set of behaviors, eliciting complex behaviors that are atypical in traditional task structures. Furthermore, the automation provided in cellworld allowed for many trials within the BotEvade task, demonstrating the effectiveness of the system for modeling ethological behaviors with the control and repeatability needed for laboratory experiments.

Discussion

In this study, we describe a system that allows researchers to study animal interactions with a robotic agent, enabling a rich set of task designs set within an arena with adjustable spatial complexity. The physical basis of the system is a modular arena which allows flexible configurations of obstacles within a 2.56 m^2 open field, supplemented with automated doors and feeders. The entire arena is monitored by a high speed tracking system, allowing the robot to react to the behavior of an animal with a lag of a hundredth of a second. Here, we leveraged this system to create a predator-prey-like task, in which we trained mice to evade a robot equipped with an aversive airpuff mechanism as it traversed a complex arena. We found that mouse behavior was strongly modulated by both the complexity of the arena and the presence of the robotic predator, finding that mice took more varied paths when compared to predator-free open fields, and observing examples of more complex behaviors, such as baiting and peeking.

While cellworld is capable of replicating previously published behavioral tasks (Fig. 2d), we argue that this system also introduces some distinct advantages over prior approaches. The two key innovations deployed here are 1), a mobile agent whose behavior is coupled to that of the experimental subject and, 2) a large, rapidly reconfigurable arena. Below, we detail how these two experimental features allow experiments that are challenging, if possible at all, using current methods.

Previous studies have utilized robotic agents to study rodent behavior, most of which fall into two main approaches: a robot that moves but is non-reactive to the animal or robots that are mostly stationary, but react when the animal comes within range. In the studies that implemented non-reactive control, the robot either moved randomly until it hit the arena wall [35, 36, 37] or was supplied with a set of predefined destinations to navigate towards [38, 39]. In the studies that implemented mostly stationary, reactive robots, the robot remained stationary until the rodent came within a specified range, after which the robot "surged" towards the mouse [40, 41, 42, 43]. Finally, most similar to the present work, there are several

studies that implemented reactive mobile robots. This includes a robot that chases the animal but is otherwise neutral [44], and a "robotic-rat" which is capable of aversive, neutral, and friendly reactions to the behavior of real rats [45, 46].

The autonomous robot within cellworld improves over these previous approaches in several respects. First, robotic control is fully reactive to the position of the mouse. This is in contrast to previous studies in which the robot did not react to the rodent at all [38, 39] or in which the rodent received foot-shocks when in the proximity of the robot, but the robot's behavior was otherwise unaffected by the animal's position [35, 36, 37]. Additionally, other studies which did use real-time sensing to react to the rodent provided very simple reaction modes, limited to a forward lunge followed by a retreat to the original position[40, 41, 42, 43]. In the current study, we improved over these previous implementations by using closed-loop control of the robot's behavior. This enabled the robot to chase the mouse with high accuracy over large distances, while still deploying aversive stimulation (airpuffs) to create negative-valence interactions.

Second, while we focused on an aversive stimulus mode in the current study, we found that without the airpuff, the robot was not inherently threatening to the mice (some possible alternatives: Supplementary Fig. 3) as indicated by previous studies [39, 44, 46]. When the airpuff was disabled, we found occasions where mice would climb onto and stay on the moving robot (Supplementary Movie 4), suggesting that the fear response was specific to the airpuff stimulus. With the airpuff module being easily removed, we can interchange the top half of the robot to any feasible mechanism as long as it does not interfere with robot navigation. This provides a wide variety of possible interaction models ranging from aversive to appetitive stimuli, in contrast to previous studies using robotic stimuli, which were only capable of inducing fear responses [40, 42, 41, 43, 35, 36, 37]. Notably, some previous studies manipulated the valence of the robot, either by baiting it with food [39] or by engaging in "friendly" biomimicry (ie. when the real rat grooms the robotic "rat" grooms) or stressful attacks [45, 46]. In line with this previous work, the system described here will be useful to study social interactions within large, complex environments. Supplementary Table 1 lists some alternative experiment paradigms including appetitive and social modes.

Finally, the robot in this study was capable of navigating a large, occluded environment, creating a two-dimensional interaction space between the mouse and the robot. Many previous studies used interaction spaces that were effectively one-dimensional, limiting the mice to a narrow corridor with the robot at one end [40, 41, 42, 43]. This resulted in very stereotyped escape and freeze behaviors that were only characterized in one of the studies mentioned [41]. By creating a large, occluded, two-dimensional interaction space, we found evidence for complex behavioral sequences of evasion and information gathering between the mouse and robot (Figures 6, 7) in addition to more stereotypical instances of thigmotactic escapes and freezing. The long sight lines and many route options through the occluded arena will be of great utility in the study of planning in the presence of a dynamic threat, which we believe to be a significant advance of the cellworld system over prior work.

Another key feature of cellworld is the

reconfigurability of the obstacles within the arena. This allows the experimenter to recreate traditional experimental setups (Fig. 2d), or create other desired layouts. Existing experiments studying memory, navigation, decision-making, and planning typically take place in an open field [47, 16, 48], or in highly simplified mazes [18, 49, 50]. Intuitively, these arenas have little in common with natural spaces, where occluded and open areas are commonly intermingled, providing locations to hide and gather information, for instance, while evading a predator. It is possible that the simple layouts and tasks commonly favored in neuroscience may alter the cognitive processing of animals behaving in these spaces compared to the natural contexts in which they evolved, which largely motivated the creation of cellworld as an alternative platform for studying behavior. As such, how spatial complexity affects behaviors and neural representations within a given space remains an underexplored question. We have demonstrated that cellworld may be used to tackle these questions by leveraging its reconfigurability during an ethologically-inspired predator evasion task.

We took a two-pronged approach to understand the impact of spatial complexity: we used a generative procedure to create random arenas with a desired level of entropy (Fig. 2c and d) and developed methods for quantifying the spatial complexity of any arbitrary arena layout. Using these methods, we found that, 1) the randomly generated arenas were more spatially complex than traditional arenas, and, 2) the complexity of the random arenas was similar to the statistics of a natural landscape (Fig. 2e, additional landscapes analyzed elsewhere[1]). However, it should be noted that while we focused on one measure of spatial complexity (network degree complexity), it is likely that this metric does not capture all of the elements of a space that might be relevant for behavior (for instance, a hairpin maze is more complex than many high entropy worlds despite having fewer routes). A promising alternative for further exploration is lacunarity, a metric used by landscape ecologists which is sensitive to the spatial scale of environmental features and can distinguish between repeating versus random occlusion arrangements [51, 52, 53]. We applied this method to the mid-entropy arena used in this study, and found that this configuration closely resembled the lacunarity profile of a natural landscape (Supplementary Fig. 1g). Furthermore, our previous work suggested that mid-entropy arenas (such as the one used in this study) had the greatest utility for planning in simulations of the BotEvade task [1, 54]. These results suggested that mice evading a predator within more natural (i.e. high complexity) spaces are more likely to use planning.

Therefore, we leveraged the features of cellworld to emulate interactions with a predatory "other" within the ethologically-inspired arena. We found that the spatial complexity of the arena, paired with a mobile threat, strongly modulated mouse behavior. In the open arena, mice reverted to thigmotactic routes to evade the robot, while in the occluded arena, mice engaged in long sequences of evasion, taking longer and more diverse paths in the presence of the robot. This suggests that in the presence of threat, low complexity configurations can lead to more stereotyped behaviors while high complexity configurations can lead to more flexible behaviors, as supported by prior computational studies [1, 54]. In addition, we observed many examples of "peeking" and "baiting", actions which are

rarely observed in conventional mazes (but see [28, 29, 31]).

For example, in a trial where the robot remained stationary, the mouse repeatedly peeked, approached, and retreated from the now immobile threat, as if in response to a violation of its internal model (Supplementary Movie 10). Based on this and our previous observations, we hypothesize that these complex behaviors arose through the implementation of an internal model to predict the location of the robot and subsequent planning of routes through the complex space to avoid it. While a large amount of future research will be required to test this hypothesis, cellworld and BotEvade now provide a lab-based method to do so.

Indeed, the instances of "baiting" and "peeking" we observe resemble previous reports of deliberative behaviors, such as vicarious trial and error (VTE [23]), which coincide with neuronal activity believed to reflect planning [50, 55, 56]. However, it is unclear whether behaviors such as "baiting" and "peeking" represent planning or simply a freezing response upon sensory contact with the threatening stimulus. As with other examples of distraction displays, such as the broken-wing display of birds [27], it is possible to interpret these results without a mechanism for planning: the mouse embarks on a route towards the goal, encounters the threatening stimulus, freezes in fear, and then reroutes to escape as the threat "looms" towards it (there is evidence for neural mechanisms supporting this interpretation of the behavior[57]). While it is unclear whether these behaviors are based on explicit plans, it is clear that the interplay between the robot and environment caused these behaviors to arise. It should be noted that the "peeks" and "baits" shown are purely from hand-picked examples. However, with the large presence of them consistent across all mice, we are confident that the cellworld system allows us to repeatably capture these complex behaviors. This provides the opportunity to define more detailed methods to identify, characterize, and assess under what conditions they emerge, again establishing future avenues to study planning during ethologically inspired tasks.

Finally, we note that the inherent discretization of the honeycomb lattice of cellworld eases synergy with computational ethology as common frameworks for reinforcement learning and partially observable Markov decision process-based planning algorithms[58, 1] are in discretized rather than continuous space. Supplementary Figure 7 shows a simulation of mouse behavior based on prior work[1, 54], showing good agreement with trials from a subset of the mice. Similar simulations are underway for comparison to the measured behavior of people performing BotEvade within a scaled cellworld in virtual reality, where the robot has been replaced with a predator avatar. This simulation environment is being readied for release along with a future publication.

Traditionally, neuroscience has favored behaviors and stimulus modes that are easily repeatable and measurable in the laboratory [18, 49, 47, 59, 16, 48, 60, 61]. More recently, advances in recording methodologies have allowed neuroscientists to record from increasingly large numbers of neurons [62, 63, 64, 65], and the rise of machine learning has provided many tools for quantifying natural behaviors [66, 67]. With these advances, there is a push to leverage behavior to study the brain [24, 25, 26], but it is unclear how neuroscientists can balance

the repeatability of traditional experimental setups with the need to elicit and quantify natural behaviors. Here we provide a solution to this problem through a modular system that allows flexible behavioral task design, closed-loop control of a mechanical agent, and detailed video monitoring. We show that we can reliably train rodents to perform hundreds of trials per hour in the presence of an aversive robot, and found that mice performed complex behaviors that are typically not observed or not quantified in prior studies. Furthermore, we provide a generative method for creating random arenas, and spatial complexity metrics to assess how similar the experimental space is to more naturalistic habitats. Combined, the features of this system represent a key step towards discovering and studying ethologically-relevant behaviors in a laboratory setting.

Limitations of the Study

There are several limitations to our approach. The speed of our robot is on average about 1/3 that of the mouse. This limitation is a combination of the increased robot size and mass needed with the aversive module, and consequent challenges with obstacle gaps that are near the width of the robot. Predators are often larger than prey and therefore can at times match or surpass the speed—if not the agility—of their quarry; the effects of this regime would be interesting to explore. In past tests with faster robots, we have seen a tendency to elicit more reactive responses such as thigmotaxis, but a more thorough investigation is needed once maneuvering and mass issues have been addressed.

There are several differences between a natural predator and our robot that could affect the mouse's behavior. While natural predators are sources of sound and odorants, these experiments feature frequent cleaning with ethanol and the presence of masking white noise to prevent the mouse from hearing the robot (confirmed by many encounters where the mouse startled to see a robot after rounding an obstacle). Additionally, the movement capabilities and search patterns of the robot used in this study were limited, comprised of a simple chasing strategy. However, programming the robot with more advanced strategies is perfectly feasible within the current system, and will merit further research.

Nonetheless, we suggest that BotEvade approximates predator-prey interactions. The use of a robot versus a natural predator is not itself necessarily a problem, since mice have no reason to think an unknown pursuing agent is anything other than a predator. One key difference between experimental encounters with the robot and an encounter with an actual predator is that real predator-prey interactions may result in injury or death. However, as we demonstrate here, the airpuff was sufficiently aversive to elicit escape behaviors on nearly every encounter (Fig. 5d). Therefore, even if we were to outfit our robot with a lethal mechanical bite, mice would rarely dwell within striking distance; therefore, for all the mouse knows, the robot *does* have a lethal bite. Based on these results, our apparatus is sufficient to elicit naturalistic evasion behaviors, just as the use of expanding black disk stimuli have been used in prior studies to study escape from "looming" stimuli [?] . While we expect the mouse is engaged in a predator-prey dynamic with the robot, it is the case that most predators of mice are likely to be faster, as we addressed above. This gap between the apparatus and natural behavior is likely to be

more important than the fact that the robot does not look like a natural threat and is absent a lethal bite.

Another aspect of natural predator prey interactions is that mice will often freeze when it knows it is in view of the predator in order to avoid being seen by the predator. Our robot pursues the mouse whether or not it has executed a freeze. However, altering this so that the robot only pursues moving mice needs only a very minor control code change. Whether this is appropriate likely varies between how much a predator's visual search relies on motion versus other factors such as contrast, which will in turn vary with the type of predator being considered. Here we stayed agnostic to this choice as some predators are less reliant on motion. We do permit the mouse to briefly peek around obstacles without being detected (Section d below).

Finally, we have not attempted to match natural scene statistics in cellworld, outside of the light spectrum. We made the robot black to contrast the otherwise white cellworld features and have landmarks on walls of the space (Fig. 3a). Future work should explore whether contrasting obstacle/wall/robot shapes or colors are important, and whether occlusion arrangements reminiscent of other types of natural landscapes (eg. denser or sparser arrangements akin to forest or desert environments, respectively) result in different behavioral strategies than the environments explored here.

Figure Legends

Fig. 1: Overview of the cellworld system. Magnetized movable obstacles break the rodent's line of sight to the robot and the robot's line of sight to the rodent and facilitate diverse rodent behaviors amidst mobile threats or opportunities. Multiple high speed cameras ensure continuous tracking, and high speed custom processing of the images ensures low latency between changes in the rodent's behavior and changes in the autonomous robot's behavior. Automated doors open and close to sequence the rodent through the rewards of the task under control of the experiment controller.

Fig. 2: The reconfigurable behavior arena. **a.** Exploded computer-aided design view, with the front walls removed for illustrative purposes. There are 331 magnetized hexagonal cells over an area of 2.56 m^2, with a long diagonal length of 2.34 m. Inset below shows magnetic attachment system. Not shown is a seamless acrylic and vinyl membrane between the obstacles and floor for cleaning and removal of odor cues. **b.** Photos of three configurations of obstacles (corresponding to 0.1, 0.5, and 0.9 entropy, see *Methods*). **c.** Top down view diagram of the obstacle configuration corresponding to the photos in panel **b**. The 0.5 arena matches the occluded condition for behavioral experiments with mice in this study. **d.** Configurations of cellworld to match some commonly used laboratory assays of learning and memory. Greyed out areas of the habitat represent areas not accessible to mice that were fully filled in with obstacles for spatial analysis. **e.** *Left*: Spatial complexity versus entropy. Line plot shows 500 repeats for each of 14 different entropy levels of cellworld, along with other

configurations. The dashed line represents the mode of spatial complexities in the natural landscape. *Right*: Illustration of the random sampling process used to select 1162 hexagonal cell-worlds of the natural landscape. The worlds are scaled such that each cell is 2 m in size, the approximate size of small herbivorous prey animal common in this habitat, for a total world size of 50 m. 162 samples that did not include any coverage (spatial complexity of 0) were removed for the calculation of the mode. The natural landscape is a binarized Google Earth image representing a 1941 m × 1139 m portion of the Okavango Delta in Botswana. The full color image and details of the natural landscape can be seen in Supplementary Figure 1. **f.** Histogram of spatial complexity of the worlds generated for the line plot in **e** including the spatial complexity of other configurations and patterns from **c** and **d**.

Fig. 3: The camera system and an autonomous interacting agent. **a.** Raw video from the four cameras. Note landmarks on top and bottom walls. **b.** Main outputs of the camera system including a summary of the mouse detection pipeline. *Left*: Stitched image processed from the four raw camera views. The robot predator is present, and the circle around it shows the attack threshold. *Middle*: Mouse detection process utilizing background subtraction and color-connected components. *Right*: Zoomed-in view of a mouse peeking around an obstacle at the predator robot from the four camera views. Were the video taken with the upper-right camera alone, the peeking behavior would not be registered and automated tracking would fail. **c.** Exploded view of the robot showing main components, with the aversive stimulus module used for the experiments described here. **d.** Image of the robot configured with a CO_2 canister for airpuff delivery. **e.** Front and side views of the CAD model of the robot. **f.** Top view of the robot in the arena. Inset shows a magenta circle used to depict the airpuff attack threshold (left) and background subtraction for tracking (right).

Fig. 4: The BotEvade task, modeled after predator-prey interactions. **a.** State flow diagram. The two processes that comprise the Main Process (black rectangle), for when mice are over 32 cm away from the robot, are "Pursue" and "Search". Below the Pursue Behavior node is an illustration of a typical pursuit scenario: the mouse is in view of the robot, and while in view, the robot will pursue. Below the Search Behavior node is an illustration of a typical search scenario: the mouse is out of view, and the robot randomly selects a cell out of view to go to (purple line and cell). When the mouse is less than 32 cm away from the robot, the "Attack" process (red rectangle) is triggered for releasing the sequence of two airpuffs. **b & c.** Experiment events shown alongside door events for the four automated doors (two at the start port, two at the end port). **d.** A single trial of rodent-robot interaction during the BotEvade task. A loud white noise generator prevents mice from hearing the position of the robot when it is out of view, and the arena is cleaned with

ethanol between subjects to remove olfactory cues. **e.** Composite image of the arena, with experiment lighting on the left and overhead lighting on the right for clarity.

Fig. 5: Learning the BotEvade task and aversive airpuff. **a.** Route pattern enforced by the BotEvade task. **b.** For experiments with obstacles, over a period of up to 22 days for 8 mice, there is a sequence of four phases: 1) one day of corridor training (CT) where a channel connects the start and end doors; 2) arena training (T), wherein mice run through the task at their own pace, with no robot present, until trials/min plateaus and the mouse runs greater than 15 trials per 30 minute session; 3) robot (R), where mice now are challenged with the robot predator until trials/min plateaus, followed by an additional 2 days of trials; 4) five days where mice experience the same conditions as the prior phase but without the robot (PR). For the R phase, we show the robot as configured with 360° vision for the shown position. In this and other typical robot locations, the obstacles provide many locations for the mice to avoid being seen by the robot. The total number of trials collected across all phases and mice is n = 6678. **c.** Trial count during arena training (T) for 8 mice (individual colored lines; average trace ± STD indicated by black dashed line and grey shading). Vertical dashed line shows the start of the plateau phase. **d.** Change in distance between the robot and mouse over 2 seconds following an attack (two sequential airpuffs, n = 178 attacks). Red/orange lines represent distance traces after individual attack events, while the grey distribution represents the 97.5th percentile of the distances when randomly sampling trajectories without attack events 19,430 times. If an individual trace fell outside of the 97.5th percentile of the random distribution after 1 s it was considered significant (red traces), otherwise it was colored in orange.

Fig. 6: Measurements of mouse and robot dynamics during BotEvade. **a.** Mouse trajectories from individual trials across experiments with obstacles and no obstacles. Color indicates the mouse's speed. For the with-obstacle cellworld experiments (top row), trajectories are shown for the plateau phase of the training (T) phase (n = 1615), the plateau phase + 2 days for robot (R) phase (n = 1248), and for 5 days of the post-robot (PR) phase (n = 2238) for n = 8 mice. For the no-obstacle cellworld experiments (bottom row), trajectories are shown for 2 days of the T (n = 182) and the R phase (n = 220) for n = 2 mice. **b.** Average number of trials per 30 minute session per mouse across each experiment phase (n = 8 mice). In this plot and all other box plots, the horizontal line is the median; box is interquartile range (IQR); whiskers are 1.5 times the IQR. Data points beyond the whiskers are denoted by circles. Two-tailed Kruskal-Wallis (KW) test: H(2) = 16.88, $p = 2.16 \times 10^{-4}$; post-hoc Dunn test: R vs. PR: $p_{\text{adj}} = 1.23 \times 10^{-4}$. Asterisks indicate significant pairwise Dunn's tests between corresponding phases. **c.** Average trajectory length per trial, per mouse in each experiment phase. Two-

tailed KW test: H(2) = 12.44, $p = 1.99 \times 10^{-3}$; post-hoc Dunn tests: R vs. PR: $p_{\text{adj}} = 1.39 \times 10^{-3}$). **d.** Example clustering results for one mouse in all of the robot experiments (left) and all of the post-robot experiments (right). Trajectories are colored by their cluster assignment, while the average trajectory for each cluster is indicated by the solid lines outlined in white. The average trajectory thickness was proportional to the number of trajectories included in the cluster. *Inset*: cluster distance was determined by averaging the distance between each individual trajectory and the closest cluster. **e.** Average number of clusters per mouse in each experiment phase. H(2) = 15.54, $p = 4.22 \times 10^{-4}$; post-hoc Dunn test: R vs. PR: $p_{\text{adj}} = 2.61 \times 10^{-4}$. **f.** Average distance from the nearest cluster per mouse in each experiment phase. H(2) = 18.02, $p = 1.22 \times 10^{-4}$; post-hoc Dunn test: R vs. PR: $p_{\text{adj}} = 6.63 \times 10^{-5}$. **g.** Average moving speed per mouse in each experiment phase. Two-tailed KW test: H(2) = 8.34, $p = 0.015$; post-hoc Dunn test: R vs. PR: $p_{\text{adj}} = 0.027$. **h.** Distribution of pause duration at the entrance in each experiment phase (colors as in b). Two-tailed KW test on pause frequency for each phase: H(2) = 12.39, $p = 2.03 \times 10^{-3}$; post-hoc Dunn tests: T vs. R: $p_{\text{adj}} = 0.016$, R vs. PR: $p_{\text{adj}} = 3.43 \times 10^{-3}$. **i.** Distribution of pause duration in the arena in each experiment phase (colors as in b). Two-tailed KW test on pause frequency for each phase: H(2) = 10.14, $p = 6.28 \times 10^{-3}$; post-hoc Dunn tests: T vs. R: $p_{\text{adj}} = 0.017$, R vs. PR: $p_{\text{adj}} = 0.017$. *Inset*: duration distribution of long pauses (>2 s). **j.** Average number of long pauses per trial in the arena. Two-tailed KW test on pause frequency for each phase: H(2) = 10.03, $p = 6.62 \times 10^{-3}$; post-hoc Dunn test: R vs. PR: $p_{\text{adj}} = 8.94 \times 10^{-3}$, T vs PR: $p_{\text{adj}} = 0.044$.

Fig. 7: Behaviors we term "baiting the robot" and "peeking". **a.** An example of "baiting behavior". Moments of "baiting" are highlighted with the blue numbered circle. The trajectory of the mouse and robot is color coded by speed. Here, (1) the mouse comes to a point where it is seen by the robot, which is typically at a random location near to the goal at the start of a trial. The mouse (2) then retreats, provoking the robot to pursue (3). This retreat-pursue cycle repeats (4–5) until the robot is close to the start gate (6), at which point the mouse uses its superior speed to outmaneuver the robot by running along the north wall. This trial is shown in Supplementary Movie 8. **b.** An example of "peeking" behavior, where the mouse (cyan dot) makes initial contact with the robot, retreats, and then appears to reconfirm the robot's location (magenta dot) by making visual contact with the robot while hiding its body. The peeking event is highlighted with the red numbered circle. The trajectory of the mouse and robot is color coded by speed. *Left*: The mouse encounters the robot, then retreats behind obstacles. *Middle*: From the concealed location, the mouse peeks and makes line of sight with the robot at 10.28s (magenta rectangle indicates the robot hull). *Right*: The mouse reroutes and escapes. The

legend indicates the speed of the two agents. **c.** Still frames of the mouse pose at the time of the peek (10.28s). *Left*: Mouse and robot locations in the experiment frame. Dashed lines and open area indicates the binocular field of the mouse (head direction $\pm 20°$), which was calculated using DeepLabCut annotations of the mouse pose [66]. *Middle*: The corresponding frame of stitched video with tracking annotations. *Right*: The zoomed-in view of the mouse during the peek, indicated by the black square in the middle panel. The mouse's stretch attend posture [29] is evident. This trial is shown in Supplementary Movie 9.

STAR Methods

Resource availability

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Dr. Malcolm MacIver (maciver@northwestern.edu)

Materials availability

Instructions to develop the specialized hardware of this study, including the cellworld arena and robot, are shown on the following website: `https://cellworld.github.io/`. Otherwise, any other requests for materials can be directed to the lead contact.

Code and data availability

- All the data used in this study has been uploaded to Github and can be accessed using the link in the resource table or from the following website: `https://cellworld.github.io/`.
- Premade software packages utilized in this study are listed in the resource table. Custom code used for data analysis can be accessed from the following website: `https://cellworld.github.io/`. Code developed for the cellworld system, such as those used for the robot, camera system, and doors, were added to Github and can be accessed using the link in the resource table.
- Any additional information required to reanalyze the data reported in this work paper is available from the Lead Contact upon request.

Experimental model and study participant details

In this study, we used a cohort of eight adult *Mus musculus* (C57BL/6, Charles Rivers Laboratories, 8-10 weeks of age at the start of experiments) mice containing four females (labeled as FMM9, FMM10, FMM13, and FMM14) and four males (labeled as MMM10, MMM11, MMM13, and MMM14). All mice were single-housed during experiments, at 28 degrees Celsius on a 12-h light:dark cycle with food provided ad libitum. All mice underwent water scheduling before the start of the training such that they were restricted to 75% of their initial weight. Initial weight was determined by taking the average weight across 3 consecutive days under normal water and food supplies. Once at the correct weight percentage, all mice ran one 30 minute experiment every weekday following the same experimental phase sequence. All experimental procedures were in accordance with NIH guidelines and approved by the Northwestern Animal Care and Use Committee.

Method details

The cellworld

Cellworld is approximately 2.34 m long at its widest section, consisting of a large open field labeled as the "arena" (Fig. 2). The arena is surrounded by 1.06 m tall walls (Fig. 2a-b). The entire structure of cellworld is made primarily of laser-cut white acrylic. The arena itself consists of 10 acrylic pieces engraved with a hexagonal grid consisting of 331 magnetic hexagon cells each roughly 11 cm apart from center to center (Fig. 2a, inset). Two 11.11 mm diameter neodymium magnets (D74-N52, K&J Magnetics, Pipersville, PA, USA) are placed into each cell. A thin 3.175 mm layer of clear acrylic followed by a 3.175 mm layer of clear vinyl cover the arena floor. Silicone sealant is applied to all corners and joining structures that the mice may interact with, containing debris and allowing for more thorough cleaning. Obstacles are 17.7 cm tall and made out of white acrylic. The base of the obstacle is the size of a cell and has two neodymium magnets that are identical in size, type, and location to the cells of the arena floor. This attracts the obstacle to the cells in the arena and lets us freely place and change the configuration of cellworld's environment (Fig. 2a).

There are two "chambers" at the start and end of the arena, and an external "mouse return chute" connects these chambers (Fig. 2a). This forms a loop where mice are introduced to the start chamber, traverse the arena, enter the end chamber, and traverse back to the start through the return chute. Water rewards are located at each chamber as motivation for the water-scheduled mice. Doors connected to a Raspberry Pi system (3B+, Raspberry Pi, Cambridge, England, UK) are placed at the entrances and exits of the chambers. Both the doors and water feeder are fully autonomous. The doors are primarily built out of laser cut white acrylic and 3D-printed polylactic acid (PLA) parts and use micro DC motors (50:1 6V micro metal gearmotor, Pololu, Las Vegas, NV, USA) in combination with limit switches to set and detect the open and closed states of the door. The base of the door that interacts with mice is made out of neoprene rubber material, which we have found to be strong enough to not be damaged by mouse manipulation but pliable enough to not harm mice if the door were to close on them. Each water dispenser is controlled by a metal lick port connected to a capacitive sensor (AT42QT1011, SparkFun, Boulder, CO, USA) and Raspberry Pi. When a mouse licks the spout, a capacitive signal is sent to the Pis, opening a solenoid valve for a fixed amount of time. The time was calibrated so that 2 µL of water was given per reward during a lick—a total of 4 µL of water is given per trial.

Lighting in cellworld was tuned to match crepuscular light conditions and provide the mice a more naturalistic environment (Fig. 4e). This involved using a combination of a LED full spectrum bulb (9-Watt LED Grow Light Bulb, General Electric, Boston, MA, USA) with a purple gel light filter attached, a LED UV bulb (UV LED Black lights Bulb, SHGPODA, Shenzhen, Guangdong, China), and lighting soft boxes for light diffusion. This lighting configuration emulated the spectrum and illuminance of real-world measurements during twilight with an energy peak around 400 nm and an overall illuminance of 2 lux [68, 69]. Two additional red lights (660 and 850 nm; LED Red light therapy bulb, Wolezek LED, China) in soft boxes were added to improve camera visibility, but are likely to be far enough

outside the visual range of mice vision to not interfere with the crepuscular lighting [70, 71].

Camera system hardware

The computer vision system employs four advanced cameras (three Basler axA2040-180km and one Sentech STC-CMB401PCL), comprising high-speed, low-latency, high-definition, and infrared CMOS sensors, which interface with the central computer via a PIXCI CL1 frame grabber PCI cards (Epix, Inc., IL, USA) using CameraLink interface. Video frames were acquired as 10-bit grayscale images at a resolution of 2040×2040 pixels at 120 fps, generating a dataset of approximately (4 cameras × 10 bit × 2040 × 2040 × 120 fps) 2.3 Gb/s. The central computer used to process the data stream from the cameras is equipped with the Ubuntu 22.04 operating system and features an Intel i9-10920X CPU, an NVIDIA GTX 3090 GPU, and 64GB of DDR4 RAM.

Unified field of view

The use of multiple cameras provides comprehensive coverage of the arena, capturing scenes that may otherwise be obscured. However, this necessitates the real-time amalgamation of individual camera feeds into a unified view with minimal latency. Moreover, the system must accurately correlate pixels to physical locations. This precision is critical to safeguard interactions with animal subjects during trials and to maintain the integrity of the experiment outcomes, given that the computer vision system is the primary data source. Any error in this translation could potentially skew results and their interpretation.

In conventional image stitching, the acquisition order or arrangement of source images is typically unspecified. Consequently, the stitching algorithm must identify the overlapping regions between individual images. This task is commonly achieved through the use of Scale-invariant feature transform (SIFT)—a technique that operates on the gradient of the raw image and generates a series of normalized Keypoint descriptors. Following the processing of all images, the lists of descriptors that exhibit the best correspondence in a given image pair can be utilized to ascertain the degree of image overlap. An undesirable side effect of this approach is that the plane into which images are merged is computed on the fly to optimize the overlapping surface, thereby precluding the determination of the correlation between a pixel and its physical location.

Achieving high-performance, pixel-accurate image stitching necessitated the development of a custom stitching process that conforms to the system's accuracy and performance requirements. The new process employs a predefined destination plane that matches a scaled version of the arena. Instead of matching features between raw images, it uses known locations which were annotated in the raw camera images, a process called homography. We then used those locations in the predefined plane and the annotated pixels corresponding to those locations to enforce their merging into the specified Field of View (FOV). In contrast to the default blending method, the cameras' locations determine the arena section best covered by each camera. These modifications enable the pre-computation and reuse of homographic information for all cameras, simplify the merging process, and ensure a high-quality match between pixels in the composite image and their physical locations. Performance testing

results indicate that the new stitching process merges images in an average time of 8.3 ms without optimizations, which reduces to 1.8 ms with CUDA optimizations. The homographic information is configured prior to the experiment execution and it is only updated if the cameras undergo any displacement or rotation. The configuration process consists of the identification of a group of five known locations from the arena in the captured images from each camera.

To verify the accuracy of the stitched, composite image, we took advantage of the physical features of the arena: all the cells in the arena have 2 magnets, separated 3 cm apart and oriented vertically. These magnets are visible in the stitched image and their location can be computed in the physical space. To verify that the perspective correction procedure during the stitching process did not distort the true image, we manually annotated the locations of the magnets for every obstacle in the uncorrected images from the cameras (indicated by the blue x in Supplementary Figure 2a). We then used the transformation calculated from the camera calibration process to generate the expected locations of the magnets based on our knowledge of their spacing and positioning. We found that the reconstruction error from the pixel plane to the laboratory coordinate frame was on average 0.15% with a max of 0.45% (Supplementary Fig. 2a–b).

Animal tracking

Prior to the introduction of the mouse to be used as background, a stitched image of the arena is captured and stored. Every time a new frame is produced by the cameras, the stored background is subtracted from the current image to eliminate all the static features. Upon completion, the resulting matrix is converted to binary by applying a threshold followed by two consecutive cycles of dilation-erosion. The result is a binary matrix containing 0 values for static background pixels and 1 for dynamic features. Next, Color Connected Components (CCC) [72] is employed to group the islets of positive values and identify all the dynamic elements present in the image, which are then individualized. This set of dynamic elements is referred to as "detection candidates". Finally, the system endeavors to match the candidates with a collection of profiles. The list of profiles is supplied as a configuration and is characterized by a lower and upper bound limit for the pixel area of the candidate to be compared against. During trials, this parameter has been fine-tuned to match different species of mice. Four distinct species were successfully tested: *Mus musculus*, *Peromyscus maniculatus*, *Peromyscus polionotus*, and *Onychomys torridus*.

We estimated the latency and throughput of the tracking system by supplying a static image for mouse and robot detection, effectively isolating processing time by removing image acquisition from the pipeline. Using standard computing hardware, we found that our system reached a throughput of 120 fps, with a latency of less than 15 ms (measured as the time elapsed from image acquisition to detection of both agents). We were able to improve performance using CUDA optimizations, reaching throughputs of 206 fps with an average latency of 3.2 ms (Supplementary Fig. 2c).

Robot tracking

To enable real-time location and orientation tracking of the robot, three LEDs were added to its top, arranged in an isosceles

triangle configuration with the shortest side situated at the back as it is shown in Figure 3f. During the experiments, every time a new frame is produced by the camera system, a brightness threshold is applied resulting in a binary matrix with values of 1 for brighter pixels and 0 for darker pixels. As done for tracking the mouse, CCC was used to identify triangular pattern signals. If a signal consistent with the specification is found, the middle point between the shortest side (back) and the opposite vertex (front) is selected as the robot center location. Then, the orientation is computed as the vector defined by this location and the front LED.

Video post-processing

For each experimental trial, cellworld produces three types of video. The primary video log presents the unified field of view, superimposes tracking markers of the robot and mouse, and incorporates all pertinent experimental information. The system also generates raw video that includes the unprocessed images from each of the four cameras. Additionally, a multi-view, subject-centered video is produced, offering a cropped perspective of the mouse as captured from all four camera angles. This subject-centered video mitigates arena geometry interference during post-processing. Finally, raw video of the unprocessed images from all four cameras is also saved. We used the multi-view, mouse-centered video data to train DeepLabCut for offline analysis of the mouse pose [66]. This allowed us to measure the gaze angle and head location for analysis of peeks (Fig. 7b, c).

Robot hardware

The predator robot utilized in the experiments was custom-built. The skid-steer drive robot is powered by the (ESP32-WROOM-32D, Espressif Systems, Shanghia, Shanghai, China) and is driven by two geared DC motors with magnetic encoders (Geared DC Motor with Magnetic Encoder Outputs - 7 VDC 1:20 Ratio, Adafruit, New York, New York, USA). The robot was equipped with three LEDs for detection and localization via the camera sensors. The robot housed two custom printed circuit boards: one provided the supporting circuitry for the microcontroller and the motor drivers (DRV8833PW, Texas Instruments, Dallas, Texas, USA), while the other powers the LEDs and the driver used for the motor component of the puff mechanism.

Robot tracking perspective correction

An unexpected issue arising from the robot tracking setup pertains to the optical perspective of the cameras. Since the cameras are affixed to the ceiling and the LED triangle is located at the height of the robot rather than at ground level, the triangle position in the captured image shifts further away from the actual robot location the further the robot is located from the center of the image. This introduces deviations of up to 3 cm in measurements, sufficient to prevent successful navigation through gaps of 9 cm given the 12 mm side clearance around the robot. To address this challenge, it was necessary to compute the real-world physical point corresponding to the center of the image captured by each camera at ground level using the previously described homography and the known height of the robots and cameras. Based on the distance from the center of the camera plane, we calculated and accounted for the perspective drift as a function of the robot's position.

Robot controller

To control the robot in the complex environment, we established a hierarchical control system comprised of three levels: behavior, path planning, and low-level motor control. During a trial, the high-level behavior controller selected robot destinations based on tracking information from the camera system and the current state of the experiment (Fig. 4a, b). Before each trial started, the robot navigated to a spawn cell (a cell not visible to the mouse, in a region of the arena furthest from the start gate) and stopped. These spawn constraints were implemented to ensure that the robot was not visible to the mouse from the start port with the gate open, which was found to lead to long delays before the start of the trial. Then, once the mouse entered the arena the *Main process* of autonomous motion began, during which the robot observed all regions in the arena that were not obstructed by obstacles relative to its current location.

Its behavior switched between aggressive pursuit or random search depending on whether the mouse was visible or hidden to the robot, respectively. The robot only entered pursuit mode if the visual ray between the mouse and robot passed outside of a buffer zone around each obstacle that was 125% the standard obstacle size or after 0.5 seconds if the ray passed within the buffer zone. This allowed the mouse to "peek" at the robot without being immediately pursued. This motion was interrupted by the *Attack process* if the mouse was within 32 cm of the robot, during which the airpuff mechanism was triggered and released two aversive airpuffs in rapid succession. To deter excessive anxiety in the mice and enable evasion post-attack, the time between attacks was regulated to be at least 0.5 seconds apart, regardless of the distance between the mouse and robot.

The middle-level path planning controller was a hybrid proportional (P) and proportional integral derivative (PID) controller that followed intermediate waypoints on trajectories created by a standard shortest-path algorithm (A* [73]) to reach the destinations the behavior controller assigned. Finally, the low-level embedded controller received attack and speed commands via Wi-Fi. It utilized encoder feedback for motor speed control and included a state machine to manage the airpuff mechanism. For the results shown in this study, the robot autonomously navigated the task environment for 1,941 trials with no human intervention.

Robot path planning controller

The tracking system provided robot state (position and orientation) feedback for the path planning controller. During the control process, the path planner selects the furthest visible cell to the robot on the robot's desired path as an intermediate target. Then, a hybrid P- and PID- controller:

$$u(t) = (\frac{1}{(a\Delta\theta(t))^2 + 1}) * P(\Delta s(t)) \pm PID(\Delta\theta(t)))$$

is used to correct along-track error Δs (distance from target) and heading error $\Delta\theta$ (difference between desired and actual heading), respectively (Supplementary Fig. 2). To avoid collisions in cluttered environments with tight spaces, the $\frac{1}{(a\Delta\theta(t))^2+1}$ term prevents the robot from translating too quickly if the $\Delta\theta$ is large, where a is an arbitrary design parameter.

Additionally, to augment obstacle avoidance, a type of potential field-based obstacle avoidance algorithm works to repel the robot from occlusions by perturbing the desired heading angle of the robot, where the

distance between the objects largely influences the magnitude of the perturbation. This algorithm is of the form

$$\vec{F}_{perturb} = \sum_{n=1}^{O} \frac{weight}{distance^{decay}} * direction$$

where O represents all obstacles in front of and within 0.35 m of the robot. *weight* and *decay* are design parameters and *distance* is the Euclidean distance between the robot and a given obstacle. The *direction* is a unit vector orthogonal to the robot's heading direction, pointing either left or right (from the robot's perspective) depending on which side of the robot the obstacle is located. Ultimately, this perturbation force works to repel the robot away from nearby obstacles by slightly offsetting the target location. This perturbation value is updated at 50 Hz; which means it is constantly adjusted based on the location of the robot in the map during the path-following process.

Aversive airpuff

To accurately simulate the predatory behavior during the BotEvade task (Results section), we added a stimulus module (Fig. 3c) to the robot's chassis. The module is comprised of a 16 g CO_2 canister and inflator (Ultraflate, Genuine Innovations, San Luis Obispo, CA, USA), an air nozzle, a brushed DC motor (120:1 Mini Plastic Gearmotor HP, Pololu, Las Vegas, Nevada, USA), a motor shaft adapter, and 3D printed PLA parts. The custom PLA parts consist of a lever arm, a cam, and a robust framework that facilitates efficient CO_2 canister replacement and simple integration with the robot.

The airpuff mechanism of the stimulus module is triggered when the camera sensors detect that a mouse has crossed the attack threshold (Fig. 4a). For each attack event, the mechanism releases two successive airpuffs, each lasting approximately 100 ms, with a 200 ms interval between them.

During each attack (sequence of two airpuffs), the motor rotates approximately 350 degrees in one direction, is halted by a mechanical stop, and then rotates the opposite direction back to its start position. Each CO_2 canister is able to produce at least 30 strong airpuffs (generally 15 attacks, assuming one attack per encounter; Supplementary Video 1), which is more than enough to complete a 30 minute session without replacement. The canister is replaced at the end of each 30 minute session to ensure consistency in puff strength across trials.

Spatial complexity metrics

A key aspect of the design of cellworld is reconfigurability guided by measures related to spatial complexity.

Cellworld entropy. The first and most basic measure of spatial complexity used in this and earlier work [1] is Shannon entropy. This is computed with the formula for the Shannon entropy [74] of a binarized version of an arena, where each open cell is 0 and each cell with an occlusion is 1. The resulting binary matrix is turned into a vector.

Entropy is determined by the following formula:

$$e = -\left(\frac{O}{C}\log_2\left(\frac{O}{C}\right) + \frac{C-O}{C}\log_2\left(\frac{C-O}{C}\right)\right)$$

where O is the number of occlusions, and C is the total number of cells in the arena.

Shannon entropy is an effective complexity measure in the context of our random generative algorithm for cellworlds, as it presupposes no interdependence between

individual elements. It is therefore insensitive to structured patterns such as the checkerboard illustration in Supplementary Figure 1: the probability of the checkerboard pattern occurring is the same as any other pattern: $1/2^{(20\times20)}$. Despite its seeming simplicity and intuitive orderliness, in the context of our generative algorithm and as measured by Shannon entropy, each cell in the checkerboard is equally likely to be occluded or open, implying the checkerboard's maximal entropy. In this case, our intuition of the entropy of the checkerboard (low because orderly) comports better with a different concept of entropy, known as *causally conditioned entropy* [75], which is not utilized in this study. Causally conditioned entropy considers the effect that the values of prior elements have on the probabilities of subsequent ones. In the context of the checkerboard example, this approach would yield a conditional entropy value of 0, as the probability of a square being occupied is determined entirely by the preceding square's state.

Occupancy. The percentage of the space with obstacles to sensory perception ($O/C \times 100$, where O and C are as defined above). Supplementary Fig. 1a plots Shannon entropy versus occupancy and maps where the various spaces we have considered fall on the curve. Similar to entropy, occupancy places no demands on where the obstacles are. This measure is similar to the informal notion of how cluttered a space is.

Network Degree Complexity. The network degree complexity [1] provides a succinct description of the uncertainty associated with the sensory connection distance between two agents in a space, based on any two randomly chosen locations within that space. For this study, we solely consider vision.

To compute the vision-based network degree complexity for a given configuration of cellworld, we translate it into a corresponding graph. Each hexagonal cell represents a graph node. An edge exists between two nodes unless an obstacle blocks the line of sight connecting the centers of the associated hexagonal cells in cellworld, implying the visual connection is disrupted. For instance, in an obstacle-free cellworld configuration, every node connects to every other node. Consequently, the degree of each node, representing the number of its connecting edges, stands at 331, which matches the total cell count in cellworld.

To determine the Network Degree Complexity for a particular cellworld configuration, we first form a vector containing the relative frequencies for each feasible degree value greater than zero, ranging from 1 to the total number of non-obstructed cells. Then, this vector is used to compute the Network Degree Entropy. Finally, we normalize this value by the system's maximum possible entropy, which is achieved when all distinct degrees have equal probabilities.

Network Degree Relative Frequency:

$$f(d) = \frac{n_d}{D}$$

Network Degree Entropy:

$$H = -\sum_{d=1}^{D} f(d)\log(f(d))$$

Network Degree Complexity:

$$C = \frac{H}{\log(\frac{1}{D})}$$

In these formulas, n_d represents the number of nodes with degree d in the graph. D is the maximum degree possible, equivalent to the node count in the graph. The term $-\log(\frac{1}{D})$ signifies the entropy when assuming a uniform frequency across all feasible

degrees ($-\sum_{d=1}^{D} 1/D \log(\frac{1}{D}) = -\log(\frac{1}{D})$). In the earlier discussed fully connected scenario, the Network Degree Complexity is zero. However, as we introduce occlusions, the system exhibits a diverse combination of high-degree (large fields of view) and low-degree (small fields of view) nodes, leading to an increase in complexity.

To compute the complexity of the natural environment presented in the main text, we utilized projections of the scaled-up cell-world arena on a satellite photo of a natural setting. We randomly placed these projections within the image. The projection size was determined based on the observed cell-animal body ratio from the physical setup and measurements of mice (body size ≈80 mm) and cell size (≈12 cm) and the impala, a prey animal native to the habitat (body size ≈130 cm) and cell size (≈2 m).

The original color satellite photo, sourced from Google Maps, was first converted into 8-bit grayscale. Subsequently, it was binarized using a midpoint threshold. This image, originally sized 8192 × 5067 pixels and representing a real-world area of 1836.77 m × 1136.1 m, was resized to 2730 × 1689 pixels. At this scale, every 3 pixels corresponds to ≈ 2 m, matching our chosen graph node scale.

For each projection, we selected a random center, ensuring it was at least 200 pixels away from any image edge to avoid overflow. The center of each cell in the arena projection was then calculated. The immediate 9 pixels (a 3×3 pixel grid, equivalent to 2×2 meters) around each cell center were inspected to check for occlusion. A cell was deemed occluded if $\lceil 0.5 \times 9 \rceil = 5$ or more of these pixels were black. We continued this process until we identified 1000 projections with at least one occluded cell.

After obtaining the 1000 non-empty projections and transforming them into cell-world configurations, we applied the same tools and methods used for analyzing the complexity of maps generated by the generative model. This ensured a consistent and equitable comparison.

Note that in our experimental work, we have manipulated the environment to minimize the contribution of other sensory modalities besides vision. For example, we use a loud white noise generator to mask the sound of the robot, and frequently clean with 70% ethanol to remove all odor cues. Nonetheless, it is worth considering the likely effect of adding sensory modalities on network degree complexity. To a first approximation, adding modalities will create additional edges in the graph. For example, imagine an owl with precise auditory localization using vision and sound to attack a rodent. Portions of the environment blocking vision will be transparent to the auditory system. This will create edges between nodes where there is no visual connectivity (and effectively reduce the number of obstacles). In an initial situation of high clutter and low complexity, the addition of audition seems likely to increase complexity. In an initial situation of medium clutter and high complexity, adding auditory perception could decrease complexity.

Lacunarity. Lacunarity (from the latin for lacuna = gap) was devised by Mandelbrot[76] (p. 310) after he observed that two fractals with identical fractal dimension could look very different. It has been applied as a multi-scale measure of spatial texture associated with patterns of dispersal on landscapes [51, 52]. Lacunarity (Λ) measures the deviation of a pattern at a given spatial scale from translational invariance [51]. If Λ is large at a given scale, then the pattern deviates a lot from trans-

lational invariance—the pattern would look different if a block of the pattern at that spatial scale were shifted to a different location; similarly, if Λ is small, then the pattern will look similar even if that block of space is shifted. Supplementary Fig. 1 provides plots of the lacunarity values of several cell-worlds and other cases, along with summary statistics for Shannon entropy, network degree complexity, and a summary statistic for lacunarity, L-value, described further below.

While chiefly used by landscape ecologists, as in our case it has also been applied for analyzing movement patterns of animals [77]. To our knowledge we are the first to apply it with specific reference to the physics of a given sensory modality, here vision, in the analysis of behavioral spaces. The gaps that we analyze using lacunarity are assumed/designed to be transparent to vision, and the obstacles between the gaps are assumed/designed to be opaque.

To compute the two-dimensional (2D) vision-based lacunarity for our samples, we take a top down view of a space and binarize the image: cells occupied by obstacles to vision are '1', and other cells are '0'. Over a set of boxes varying in size, we compute the ratio of the variance to the squared mean of the sum of the elements within the box. Because lacunarity is usually plotted on a ln–ln scale, one is added to this ratio so that $\ln(\Lambda)$ goes to zero as Λ goes to zero, giving $\Lambda(r) = \frac{\mathrm{Var}(S)}{\mathbb{E}[S]^2} + 1$, where r is the box size and S is the occupied sites by the variable of interest—visually occlusive objects in our case. The lacunarity curves that arise (Supplemental Figure 1b) gives information as to what spatial scale a given landscape transitions from being inhomogeneous to homogeneous, where homogeneous means $\ln(\Lambda) \approx 0$, and that the space would be invariant to the corresponding box size of space being translated to another location.

For example, for the checkerboard pattern of Supplemental Figure 1, the space is inhomogeneous up to the scale where the pattern repeats (at a box size encompassing 2 x 2 squares, 183 x 183 pixels for our image); after that, the space is homogeneous. The lacunarity curve therefore transitions from $\ln(1/P)$ at the smallest box size, where P is the percentage of the cells occupied (or 50% in this case, $\ln(\Lambda) = \ln(2) \approx 0.69$) to close to zero the size pattern repetition ($\ln(\Lambda) = \ln(183) \approx 5.2$). What lacunarity compactly communicates is how sparse the space is at the finest scale of analysis (the curve starts at $\ln(1/P)$), and the evolution of the curve as the box size increases to the full extent being analyzed (and therefore mathematically the lacunarity must be unity so the $\ln(\Lambda)$ plot goes to zero). Between these two limits, the descent of the curve shows the spatial scale where the pattern of the space repeats, and how quickly that transition occurs. For self-similar patterns, the lacunarity curve is a straight line on a log-log plot, with a slope equal to the fractal dimension minus the Euclidean dimension. Our natural landscape sample has a near straight line slope, and its fractal dimension is ≈1.7 (intercept ≈2.6: $y \approx (1.7 - 2)x + 2.6$). Other landscape samples, and a survey of the lacunarity values found in different types of aquatic and terrestrial biomes, are provided in earlier work[1].

The integral of the lacunarity curve, the L-value [53], provides a quick index into the magnitude of the heterogeneous space. For two spaces with similar occupancy (and thus starting near the same value of Λ at the smallest box size), if the space tran-

sitions quickly to invariance under translation, then the L-value will be small; if the transition occurs at larger spatial scales then the L-value will be large.

For example, consider the lacunarity curves for the natural landscape sample and the 0.5 Random configuration used in our experiments, Supplementary Figure 1c. As shown in the legend, the L-value for the natural landscape (8.1) is larger than Random 0.5 (7.3), with comparable initial Λ values, since the natural landscape has a shallow straight-line-like decline whereas the Random cellworld declines more rapidly. These two L-value are close, and so is the corresponding network degree complexity of the two cases. In contrast, the the L-value of the checkerboard pattern and hairpin maze are similar despite very different network degree complexities (checkerboard at 0.0 and hairpin at 0.72), grouping them into the same relatively homogeneous space category.

Limitations: There are several limitations to the sensory oriented lacunarity analysis as presented here. One is the assumption that the profile of an obstacle from above properly represents how vision interacts with the object over its height. While true by design for the obstacles in cellworld, this is not generally the case for natural obstacles as trees with their narrow bases and wide tops. Some of the limitations of performing a lacunarity analysis of 3D landscapes using 2D projections can be circumvented by computing 3D lacunarity [78], but 3D scans of space are rarely available. Further, the metric has its roots in computational geometry and landscape analysis, and the application to analyzing how a landscape is sensed and processed by an animal is challenged by the difficulty of deciding on the relevant spatial scales, and the multiplicity of ways a landscape is sensed. Finally, the relevant perspective on the space for the calculation is not always obvious. It can be argued that for an application where spatial complexity is being examined through the lens of cognitive map formation, a top-down perspective such as used here may be appropriate [1]; for other forms of spatial processing, other options may be considered.

Summary. To recap, we have discussed the use of cellworld entropy, occupancy, network degree complexity, and lacunarity. Each of these quantities has different roles. The **Shannon entropy** is a practical measure that serves as a target in our generative model to produce cellworlds. The **occupancy** is easily understood as something akin to how cluttered an environment is, and also gives us one point on the lacunarity curve for the space, as it will be $\ln(\Lambda) = \ln(1/P)$ at the smallest spatial scale (termed grain), where P is occupancy. **Network degree complexity** tells us how uncertain the distance of sensory connection will be for any two randomly chosen locations within the corresponding cellworld. But high uncertainty can arise within a relatively homogeneous space as well, as it does for the hairpin maze. Finally, **lacunarity** gives us a multi-scale view of the invariance of a pattern to translation across spatial scales of interest. If you've been in a space where you feel it looks the same in all directions, and the same when you move to a different location, then at that spatial scale, the $\ln(\Lambda)$ value of the space is nearing zero. One could speculate that animals that use cognitive maps will be challenged, and need external landmarks to navigate succesfully in such spaces. The area under the lacunarity curve, or L-value, is useful when a single value to represent a space's lacunar-

ity is desired and can help group spaces with different network degree complexity values but similar levels of spatial homogeneity.

Arena configurations

Leveraging the flexibility of cellworld, the system can emulate a wide range of established experimental designs. Furthermore, cellworld can effectively reproduce environments that have ethological relevance with varying levels of visibility.

The creation of these diverse environments is realized through the use of automated tools specifically developed for this system. The model for generating these environments relies on one primary variable: the target entropy level of the cellworld. This variable can be manipulated through two parameters which include the number of occlusions within the arena, or the desired level of entropy. The method starts with zero occlusions ($O = 0$) and incrementally adds occlusions until the specified entropy level is achieved. Conversely, if the number of occlusions is explicitly given, this step can be bypassed. Finally, the process selects O cells randomly and marks them as occluded.

To guarantee reproducibility, the algorithm accepts an optional *seed* parameter. When provided, this parameter secures a consistent occlusion configuration across runs. However, in the absence of this parameter, the procedure will generate a unique occlusion configuration for each execution.

Validation Criteria. A configuration of obstacles is considered valid when it meets two essential conditions: First, all non-occluded or open cells should be connected by an open path, ensuring no open cells are left in isolation. Second, the cells representing the entry and exit points must be open, and at least one viable path between the entry and exit cells exists. The configuration generation process repeats until the resulting arrangement passes the validation criteria.

Mouse experiments

Experimental conditions were determined by one of four sequential experimental phases—corridor training (CT), arena training (T), robot (R), and post-robot (PR) phase—which the mice were assigned based on a combination of standardized and individualized progress quotas. All mice followed the same phase sequence. During corridor training, a channel made from modified vinyl gutters was placed across the length of the arena connecting the start and end doors entering and exiting the arena. The CT phase lasted for only one day, but two mice (FMM9 and FMM10) were given a second day in the corridor due to the lack of trials. The corridor was removed and obstacles were introduced to match a specified mid-entropy configuration (named as "21_05"; Random 0.5 in Fig. 2c) for the T phase. Mice roamed freely and progressed at their own rate. Once trial count was $\geq$ 15 over a 30 minute session and the trial count plateaued, they transitioned to the R phase where they performed the same task in the presence of the robot. Progressing past this phase also required a plateau of their trial count; however, we added an additional two days of experiments with the robot after this stabilization occurred. Stabilization was determined when, across a three-day window, the trial count each day did not exceed more than 20% of the three-day trial count mean. As a result, mice needed to run a minimum of three days in the T phase before we could determine if a plateau was occurring and mice could

progress. In addition, during the R phase, mice always had five days (three used for the plateau check plus two additional days) where they were considered "acclimated" to the robot to match the required number of days for the final phase (PR) where the robot was not present as a control.

The cohort of mice which experienced the no-obstacle condition (n=2, Fig. 6) underwent slightly modified training procedures consisting of previous exposure to a mid-range entropy world with and without the robot before being exposed to the open-field arena. Experiments in the open-field arena consisted of two days of self motivated exploration without the robot and then two days with the robot.

Between each mouse, all obstacles were removed from the arena, and the arena was fully wiped down using 100% ethanol. At the end of each day, the arena, the six inches of the wall closest to the ground, and all the obstacles were wiped down with 100% ethanol. The return chute and chambers were also cleaned using damp paper towel sprayed with Labsan C-Dox. During experiments, white noise was played from a white noise generator (LectroFan Classic, Campbell, CA, USA) along a nearby wall at max volume settings. Both the cleaning and white noise were applied to limit possible confounding effects from other sensory modalities such as auditory and olfactory cues.

BotEvade task

Cellworld automation is dictated by a centralized script labeled as "experiment controller", which is able to receive and broadcast experimental events from any device connected to it. In this case, the camera system, robot, and both chambers are connected to this system. As an example of how the experiment controller connects devices when a mouse licks the lick port in the start chamber, the main experiment controller receives a message from the start chamber's Raspberry Pi. Subsequently, the experiment controller broadcasts the "start trial" event to all connected components, signaling the camera system to begin saving video recordings, and for all doors to initialize to the proper state.

BotEvade utilizes a sequence of these specified experiment events to dictate the logic and progression of the task (Fig. 4a–c). To begin the task, a researcher will place a mouse in the start chamber and manually send out the "start experiment" event via a terminal connected to the experiment controller. Using this terminal, any command can be manually sent to override or alter the progression of the task. In response to the experiment starting, all doors in cellworld will close, keeping the mouse contained in the start chamber. After the "start trial" event is sent, the Raspberry Pi of the start chamber waits until the robot has reached its spawn location before the door connecting the start chamber to the arena opens. Once the mouse past a 12.7 cm radius (one cell) from the start as detected by the camera system, a "prey enters arena" event is triggered where the door behind the mouse closes and the robot begins moving. A "finish trial" event is then broadcasted once the mouse traverses the arena and reaches the lick port at the end chamber, closing/opening doors to guide it into the return chute. Simultaneously, the camera system stops recording and saves the video recording of the trial, and the robot begins to move to a spawn point for the next trial. The task will restart via the broadcast of the "start trial" once the mouse reaches the lick port in the start chamber unless the

experiment has progressed past 30 minutes. In this case, the "finish experiment" event is sent and all the doors in the start chamber will close for mouse extraction.

Behavioral analysis

We performed all behavioral analysis in Python using a custom-built library (`https://pypi.org/project/cellworld/`). Trials with tracking errors were automatically detected and removed, while a small number of trials with robot malfunctions or experimenter intervention were removed through manual inspection of the video logs. The remaining trials were counted towards the trial count for each 30 minute experiment.

Mouse and robot positions from online tracking were logged and used to perform all behavioral analyses. From the positional coordinates of the logs, we computed instantaneous speed by calculating the distance over time between adjacent frames, then smoothed the speed trace with a moving average of 10 frames (11 ms). Path length was computed as the sum of the distance between adjacent frames. To assess fear responses, we calculated the change in the distance between the robot and the mouse in a window starting at the time of the airpuff and ending after 2 seconds. To determine whether any given distance trajectory was significantly different than expected by chance, we employed a permutation procedure, where we randomly sampled n = 19,340 time points across non-airpuff trajectories and computed the mouse-robot distance. We then computed the 97.5th percentiles of the random samples, and any true distance trajectory which fell above these percentiles after 1 s post-attack was considered significant.

To determine whether mice varied their paths throughout various stages of the task, we used a clustering algorithm (QuickBundles[34]) to identify stereotyped path choices. First, we interpolated the mouse path locations along the x-axis of the arena into 100 segments of equal lengths. The interpolated paths were then clustered using QuickBundles with the following parameters: minimum number of clusters = 1, distance threshold from cluster centroid = 23.4 cm, and minimum number of paths in a cluster = 10% of total trajectories being considered. During the clustering process, the distance of each interpolated path is compared to each existing cluster's centroids to identify the minimum distance. If the distance to the closest cluster is less than the distance threshold, the path is added to that cluster and updates the cluster's centroid. If the distance exceeds the threshold, a new cluster is created for that path. Clusters that contain fewer paths than the minimum allowed (10% of trajectories in analysis set) are discarded at the end of the process, and those paths are considered unclustered. For this analysis, we pooled all paths for each mouse within each phase of the task (T, R, and PR phases) and clustered them separately. We then quantified path diversity in each phase by considering the number of clusters and the average distance to the nearest cluster for each mouse.

To detect when the mouse paused, we developed a simple algorithm which required two parameters: a distance threshold and pause duration. A pause was defined as the frames where the mouse's location remains within a given radius (distance threshold), for a given number of frames (pause duration parameter). For this analysis, the distance threshold was set to a radius of 2.5 cm and pause duration was set to 0.5 seconds.

Airpuff aversion control experiment

For analysis of fear response to the airpuff stimuli on a moving robot, we ran two naive mice in an independent control experiment where the airpuff on the robot was disabled. Both mice went through three experimental phases within an open-field arena with no obstacles. The arena doors remained closed, containing the mice inside the arena, and the mice were allowed to freely roam the arena for the duration of each daily 30 minute experiment. The cleaning procedure and white noise were identical to the main cohort experiments. The first phase served to acclimate the naive mouse to the robot and the open field environment. The robot was stationary throughout the entirety of this phase which only lasted for one day. In the next two sessions, the robot began to move and pursue the mouse as in previous experiments, but the robot's puffing mechanism was disabled. For the last phase, the robot's puffing mechanism was enabled again while the robot pursued the mouse. The no puff phase lasted for 2 days while the puffing phase lasted for 1 day.

We quantified the change in response between the puff disabled and puff enabled session by measuring the distance between the robot and the mouse 1-2 seconds after the mouse entered the attack threshold (32 cm). These data were then averaged over that time window for each puff event to calculate statistical significance (Supplementary Fig. 3b).

Quantification and statistical analysis

To calculate significance across the different experimental phases we first averaged each statistic per mouse per experimental phase. Due to the small sample size (n=8) and skewed distribution of many of the outcome measures, we used non-parametric Kruskal-Wallis (KW) tests followed by post-hoc Dunn's tests between experimental conditions for each mouse. In all post-hoc tests, the p-values were adjusted using Bonferonni correction (p_{adj}). To calculate significant differences between puff and no-puff conditions in the control experiment (Supplementary Fig. 3), we used a Wilcoxon ranksum test on the distributions distance from the robot at the time of the puff after puff-disabled and puff-enabled attacks. The statistical details of all experiments are reported in the Results section or in the legend of the associated figure, where appropriate. To indicate the results of statistical tests in figure panels, asterisks indicate the following significance levels: *p<0.05, **p<0.01, ***p<0.001, ****p<0.0001. No significance is indicated by n.s.

To account for potential gender differences we performed non-parametric Wilcoxon rank-sum tests to determine whether male and female mice differed on the task performance metrics considered in Figure 6. We found that gender did not significantly affect the number of trials performed ($p = 0.39$), path length ($p = 0.08$), average moving speed ($p = 0.56$), or the number of pauses per trial while in the arena ($p = 0.56$).

Supplementary Movies

- SM1: Movie of the aversive airpuff sequence, termed "attack" event, delivered by the airpuff module. Green timer in background to illustrate event duration.

- SM2: Movie of FMM13 fleeing from the robot following an "attack" event. Magenta circle indicates the robot's attack threshold. Magenta dot and arrow in-

dicate the robot's location and heading, respectively. When the magenta circle switches to red this indicates that an attack was triggered. Cyan arc indicates the entrance threshold. Upon initial cross of the entrance the "prey enters arena" experiment event is triggered and the robot begins autonomous motion. Cyan dot indicates the mouse's location.

- SM3: Movie of FMM16 during pilot study where mice first interacted with a pursuing robot with airpuff disabled for two sessions. This movie shows the mouse's first experimental session with the puff enabled.

- SM4: Movie of a trial with the airpuff disabled. Here the mouse climbed onto and stayed on the moving robot.

- SM5: Movie showing an example of a stereotypical trajectory without the robot (T phase) and in the presence (first clip) and absence of obstacles (second clip).

- SM6: Movie showing an example trajectory with the robot (R phase) in the presence and absence of obstacles. *First clip*: Movie of FMM10 during its seventh day of the R phase. The mouse's first line of sight with the robot occurs at frame 590. *Second clip*: Movie of trial in the open field arena with the robot (R phase). The robot shown is not the robot used during the BotEvade task, it is an earlier iteration of the robot.

- SM7: Movie of FMM9 during its first experiment session after the robot was removed (PR phase).

- SM8: Movie showing an example baiting sequence. Following the sequence described in Fig. 7a, (1, frames 1–329) the mouse comes to a point where it is seen by the robot. (2, frames 330–462) The mouse then retreats, (3, frames 463–801) provoking the robot to pursue. (4—5, frames 802–1006) This retreat-pursue cycle repeats until (6, frames 1007–1190) the robot is close to the start gate, allowing the mouse to outmaneuver the robot.

- SM9: Movie showing an example peeking event. Following the sequence described in Fig. 7b, (1, frames 1–525) the mouse makes initial contact with the robot, retreats and then (2, frames 526-579) peeks before (3, frames 580-576) rerouting and escaping the robot.

- SM10: To determine the effect on the mouse of suddenly changing the robot's behavior, this is a movie of a trial where the robot was turned off in the middle. In this trial, the mouse engages an extended sequence of peeks and reroutes as it surveys the now stationary robot.

Acknowledgements

We thank H. Davoudi for help with setting up animal care and pilot experiments. We thank L. Browdy, P. Ryan, J. Murciano, B. Zitzewitz, S. Griswold., and A. Ulmer for assisting with behavioral experiments and animal care. This work was funded by NSF IIS 2123725 and NSF ECCS 1835389 to M. MacIver and D. Dombeck.

Author contributions

M.A.M. and D.A.D. conceptualized the ideas of this study. G.E., A.T.L., and G.E.W. developed the cellworld system.

A.T.L., G.E.W., and C.F.A ran experiments for collecting behavioral data. All authors contributed to the analysis and interpretation of results and with the writing and editing of this study.

Declaration of interest

The authors declare that they have no conflict of interest.

References

1. Mugan, U. and MacIver, M.A. (2020). Spatial planning with long visual range benefits escape from visual predators in complex naturalistic environments. Nat. Commun. *11*, 3057. Number: 1 Publisher: Nature Publishing Group, https://doi.org/10.1038/s41467-020-16102-1.
2. MacIver, M.A. and Finlay, B.L. (2022). The neuroecology of the water-to-land transition and the evolution of the vertebrate brain. philtransb *377*, 1–26. Publisher: Royal Society, https://doi.org/10.1098/rstb.2020.0523.
3. Wood, R.A., Bauza, M., Krupic, J., Burton, S., Delekate, A., Chan, D., and O'Keefe, J. (2018). The honeycomb maze provides a novel test to study hippocampal-dependent spatial navigation. Nature *554*, 102–105.
4. Hoshino, S., Takahashi, R., Mieno, K., Tamatsu, Y., Azechi, H., Ide, K., and Takahashi, S. (2020). The Reconfigurable Maze Provides Flexible, Scalable, Reproducible, and Repeatable Tests. iScience *23*, 100787. https://doi.org/10.1016/j.isci.2019.100787.
5. Vallianatou, C.A., Alonso, A., Aleman, A.Z., Genzel, L., and Stella, F. (2021). Learning-Induced Shifts in Mice Navigational Strategies Are Unveiled by a Minimal Behavioral Model of Spatial Exploration. eNeuro *8*, ENEURO.0553–20.2021. https://doi.org/10.1523/ENEURO.0553-20.2021.
6. Widloski, J. and Foster, D.J. (2022). Flexible rerouting of hippocampal replay sequences around changing barriers in the absence of global place field remapping. Neuron *110*, 1547–1558.e8. https://doi.org/10.1016/j.neuron.2022.02.002.
7. Taube, J.S., Muller, R.U., and Ranck, J.B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience *10*, 420–435. https://doi.org/10.1523/JNEUROSCI.10-02-00420.1990.
8. Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature *484*, 62–68. https://doi.org/10.1038/nature10918.
9. Wood, E.R., Dudchenko, P.A., Robitsek, R.J., and Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron *27*, 623–633. https://doi.org/10.1016/s0896-6273(00)00071-4.
10. Frank, L.M., Brown, E.N., and Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. Neuron *27*, 169–178. https://doi.org/10.1016/s0896-6273(00)00018-0.
11. Howe, M.W., Tierney, P.L., Sandberg, S.G., Phillips, P.E., and Graybiel, A.M. (2013). Prolonged Dopamine Signalling in Striatum Signals Proximity and Value of Distant Rewards. Nature *500*, 575–579. https://doi.org/10.1038/nature12475.
12. Rao, R.P., von Heimendahl, M., Bahr, V., and Brecht, M. (2019). Neuronal Responses to Conspecifics in the Ventral CA1. Cell Reports *27*, 3460–3472.e3. https://doi.org/10.1016/j.celrep.2019.05.081.
13. Heimendahl, M.v., Rao, R.P., and Brecht, M. (2012). Weak and Nondiscriminative Responses to Conspecifics in the Rat Hippocampus. J. Neurosci. *32*, 2129–2141. Publisher: Society for Neuroscience Section: Articles, https://doi.org/10.1523/JNEUROSCI.3812-11.2012.
14. Omer, D.B., Maimon, S.R., Las, L., and Ulanovsky, N. (2018). Social place-cells in the bat hippocampus. Science *359*, 218–224.
15. Danjo, T., Toyoizumi, T., and Fujisawa, S. (2018). Spatial representations of self and other in the hippocampus. Science *359*, 213–218.
16. Muller, R., Kubie, J., and Ranck, J. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. J.

Neurosci. *7*, 1935–1950. https://doi.org/10.1523/JNEUROSCI.07-07-01935.1987.

17. O'Keefe, J. and Recce, M.L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus *3*, 317–330. https://doi.org/10.1002/hipo.450030307.
18. Tolman, E.C. (1925). Purpose and cognition: the determiners of animal learning. Psychological Review *32*, 285–297. Place: US Publisher: Psychological Review Company, https://doi.org/10.1037/h0072784.
19. Pfeiffer, B.E. and Foster, D.J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. Nature .
20. Gillespie, A.K., Astudillo Maya, D.A., Denovellis, E.L., Liu, D.F., Kastner, D.B., Coulter, M.E., Roumis, D.K., Eden, U.T., and Frank, L.M. (2021). Hippocampal replay reflects specific past experiences rather than a plan for subsequent choice. Neuron *109*, 3149–3163.e6. https://doi.org/10.1016/j.neuron.2021.07.029.
21. Wikenheiser, A.M. and Redish, A.D. (2015). Hippocampal theta sequences reflect current goals. Nat. Neurosci. *18*, 289.
22. Diba, K. (2021). Hippocampal sharp-wave ripples in cognitive map maintenance versus episodic simulation. Neuron *109*, 3071–3074. https://doi.org/10.1016/j.neuron.2021.09.006.
23. Redish, A.D. (2016). Vicarious trial and error. Nature Reviews Neuroscience *17*, 147–159.
24. Krakauer, J.W., Ghazanfar, A.A., Gomez-Marin, A., MacIver, M.A., and Poeppel, D. (2017). Neuroscience Needs Behavior: Correcting a Reductionist Bias. Neuron *93*, 480–490.
25. Hunt, L., Daw, N., Kaanders, P., MacIver, M., Mugan, U., Procyk, E., Redish, A., Russo, E., Scholl, J., Stachenfeld, K., et al. (2021). Formalising planning and information search in naturalistic decision-making. Nat. Neurosci. , xx.
26. Mattar, M.G. and Lengyel, M. (2022). Planning in the brain. Neuron *110*, 914–934. https://doi.org/10.1016/j.neuron.2021.12.018.
27. de Framond, L., Brumm, H., Thompson, W., Drabing, S.M., and Francis, C.D. (2022). The broken-wing display across birds and the conditions for its evolution. PROCEEDINGS OF THE ROYAL SOCIETY B-BIOLOGICAL SCIENCES *289*, 20220058. Num Pages: 10 Place: London Publisher: Royal Soc Web of Science ID: WOS:000790306100010, https://doi.org/10.1098/rspb.2022.0058.
28. Blanchard, R.J. and Blanchard, D.C. (1989). Antipredator defensive behaviors in a visible burrow system. Journal of Comparative Psychology (Washington, D.C.: 1983) *103*, 70–82. https://doi.org/10.1037/0735-7036.103.1.70.
29. Blanchard, R., Parmigiani, S., Bjornson, C., Masuda, C., Weiss, S., and Blanchard, D. (1995). Antipredator Behavior of Swiss-Webster Mice in a Visible Burrow System. Aggresive Behav. *21*, 123–136. Num Pages: 14 Place: New York Publisher: Wiley-Liss Web of Science ID: WOS:A1995RM38900004, https://doi.org/10.1002/1098-2337(1995)21:2<123::AID-AB2480210205>3.0.CO;2-E.
30. Fonio, E., Benjamini, Y., and Golani, I. (2009). Freedom of movement and the stability of its unfolding in free exploration of mice. Proc. Natl. Acad. Sci. U.S.A *106*, 21335–21340. https://doi.org/10.1073/pnas.0812513106.
31. Reinhold, A.S., Sanguinetti-Scheck, J.I., Hartmann, K., and Brecht, M. (2019). Behavioral and neural correlates of hide-and-seek in rats. Science *365*, 1180–1183.
32. Dombeck, D.A., Khabbaz, A.N., Collman, F., Adelman, T.L., and Tank, D.W. (2007). Imaging large-scale neural activity with cellular resolution in awake, mobile mice. Neuron *56*, 43–57.
33. Fink, A.J., Axel, R., and Schoonover, C.E. (2019). A virtual burrow assay for head-fixed mice measures habituation, discrimination, exploration and avoidance without training. eLife *8*.
34. Garyfallidis, E., Brett, M., Correia, M.M., Williams, G.B., and Nimmo-Smith, I. (2012). Quickbundles, a method for tractography simplification. Frontiers in neuroscience *6*, 175.
35. Telensky, P., Svoboda, J., Blahna, K., Bureš, J., Kubik, S., and Stuchlik, A. (2011). Functional inactivation of the rat hippocampus disrupts avoidance of a moving object. Proceedings of the National Academy of Sciences *108*, 5414–5418. Publisher: Proceedings of the National Academy of Sciences, https://doi.org/10.1073/pnas.1102525108.
36. Svoboda, J., Lobellová, V., Popelíková, A.,

Ahuja, N., Kelemen, E., and Stuchlík, A. (2017). Transient inactivation of the anterior cingulate cortex in rats disrupts avoidance of a dynamic object. Neurobiology of Learning and Memory *139*, 144–148. https://doi.org/10.1016/j.nlm.2017.01.003.

37. Ahuja, N., Lobellová, V., Stuchlík, A., and Kelemen, E. (2020). Navigation in a Space With Moving Objects: Rats Can Avoid Specific Locations Defined With Respect to a Moving Robot. Frontiers in Behavioral Neuroscience *14*.
38. Gianelli, S., Harland, B., and Fellous, J.M. (2018). A new rat-compatible robotic framework for spatial navigation behavioral experiments. Journal of Neuroscience Methods *294*, 40–50. https://doi.org/10.1016/j.jneumeth.2017.10.021.
39. Bos, J.J., Vinck, M., Marchesi, P., Keestra, A., Van Mourik-Donga, L.A., Jackson, J.C., Verschure, P.F., and Pennartz, C.M. (2019). Multiplexing of Information about Self and Others in Hippocampal Ensembles. Cell Reports *29*, 3859–3871.e6. https://doi.org/10.1016/j.celrep.2019.11.057.
40. Choi, J.S. and Kim, J.J. (2010). Amygdala regulates risk of predation in rats foraging in a dynamic fear environment. P Natl. Acad. Sci. Usa. *107*, 21773–21777.
41. Walters, C.J., Jubran, J., Sheehan, A., Erickson, M.T., and Redish, A.D. (2019). Avoid-approach conflict behaviors differentially affected by anxiolytics: implications for a computational model of risky decision-making. Psychopharmacology *236*, 2513–2525. https://doi.org/10.1007/s00213-019-05197-0.
42. Kim, E.J., Park, M., Kong, M.S., Park, S.G., Cho, J., and Kim, J.J. (2015). Alterations of hippocampal place cells in foraging rats facing a "predatory" threat. Current biology: CB *25*, 1362–1367. https://doi.org/10.1016/j.cub.2015.03.048.
43. Kong, M.S., Kim, E.J., Park, S., Zweifel, L.S., Huh, Y., Cho, J., and Kim, J.J. (2021). 'Fearful-place' coding in the amygdala-hippocampal network. eLife *10*, e72040. https://doi.org/10.7554/eLife.72040.
44. del Angel Ortiz, R., Contreras, C.M., Gutiérrez-Garcia, A.G., and González, M.F.M. (2016). Social Interaction Test between a Rat and a Robot: A Pilot Study. International Journal of Advanced Robotic Systems *13*, 4. Publisher: SAGE Publications, https://doi.org/10.5772/62015.
45. Shi, Q., Ishii, H., Konno, S., Kinoshita, S., and Takanishi, A. (2012). Image processing and behavior planning for robot-rat interaction. In 2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob) (Rome, Italy: IEEE), pp. 967–973. URL http://ieeexplore.ieee.org/document/6290292/.
46. Shi, Q., Ishii, H., Kinoshita, S., Takanishi, A., Okabayashi, S., Iida, N., Kimura, H., and Shibata, S. (2013). Modulation of rat behaviour by using a rat-like robot. Bioinspiration & Biomimetics *8*, 046002. Publisher: IOP Publishing, https://doi.org/10.1088/1748-3182/8/4/046002.
47. 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.
48. Hafting, T., Fyhn, M., Molden, S., Moser, M.B., and Moser, E.I. (2005). Microstructure of a spatial map in the entorhinal cortex. Nature *436*, 801.
49. O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. Experimental Neurology *51*, 78–109. https://doi.org/10.1016/0014-4886(76)90055-8.
50. Johnson, A. and Redish, A.D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. Journal of Neuroscience *27*, 12176–12189.
51. Plotnick, R.E., Gardner, R.H., and O'Neill, R.V. (1993). Lacunarity indices as measures of landscape texture. Landscape Ecology *8*, 201–211.
52. Plotnick, R.E., Gardner, R.H., Hargrove, W.W., Prestegaard, K., and Perlmutter, M. (1996). Lacunarity analysis: a general technique for the analysis of spatial patterns. Physical Review E *53*, 5461.
53. Hoechstetter, S., Walz, U., and Thinh, N.X. (2011). Adapting lacunarity techniques for gradient-based analyses of landscape surfaces. Ecological Complexity *8*, 229–238.
54. Espinosa, G., Wink, G.E., Lai, A.T., Dombeck, D.A., and MacIver, M.A. (2022). Achieving mouse-level strategic evasion performance using real-time computational planning.

55. Papale, A.E., Zielinski, M.C., Frank, L.M., Jadhav, S.P., and Redish, A.D. (2016). Interplay between hippocampal sharp-wave-ripple events and vicarious trial and error behaviors in decision making. Neuron *92*, 975–982.
56. Schmidt, B., Duin, A.A., and Redish, A.D. (2019). Disrupting the medial prefrontal cortex alters hippocampal sequences during deliberative decision making. J. Neurophysiol. *121*, 1981–2000. Publisher: American Physiological Society, https://doi.org/10.1152/jn.00793.2018.
57. Evans, D.A., Stempel, A.V., Vale, R., Ruehle, S., Lefler, Y., and Branco, T. (2018). A synaptic threshold mechanism for computing escape decisions. Nature *558*, 590–594. Number: 7711 Publisher: Nature Publishing Group, https://doi.org/10.1038/s41586-018-0244-6.
58. Silver, D. and Veness, J. (2010). Monte-Carlo planning in large POMDPs. In nips. pp. 2164–2172.
59. McNaughton, B.L., Barnes, C.A., and O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. Experimental Brain Research *52*, 41–49. https://doi.org/10.1007/BF00237147.
60. Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsáki, G. (2008). Internally Generated Cell Assembly Sequences in the Rat Hippocampus. Science *321*, 1322–1327. Publisher: American Association for the Advancement of Science, https://doi.org/10.1126/science.1159775.
61. Jadhav, S.P., Kemere, C., German, P.W., and Frank, L.M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. Science (New York, N.Y.) *336*, 1454–1458. https://doi.org/10.1126/science.1217230.
62. Jun, J.J., Steinmetz, N.A., Siegle, J.H., Denman, D.J., bauza, m., barbarits, b., Lee, A.K., Anastassiou, C.A., Andrei, A., Aydın, C., et al. (2017). Fully integrated silicon probes for high-density recording of neural activity. Nature Publishing Group *551*. https://doi.org/10.1038/nature24636.
63. Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C.B., Carandini, M., and Harris, K.D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. Science *364*, eaav7893. Publisher: American Association for the Advancement of Science, https://doi.org/10.1126/science.aav7893.
64. Steinmetz, N.A., Aydin, C., Lebedeva, A., Okun, M., Pachitariu, M., Bauza, M., Beau, M., Bhagat, J., Böhm, C., Broux, M., et al. (2021). Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. Science *372*, eabf4588. Publisher: American Association for the Advancement of Science, https://doi.org/10.1126/science.abf4588.
65. Zong, W., Obenhaus, H.A., Skytøen, E.R., Eneqvist, H., de Jong, N.L., Vale, R., Jorge, M.R., Moser, M.B., and Moser, E.I. (2022). Large-scale two-photon calcium imaging in freely moving mice. Cell *185*, 1240–1256.e30. https://doi.org/10.1016/j.cell.2022.02.017.
66. Mathis, A., Mamidanna, P., Cury, K.M., Abe, T., Murthy, V.N., Mathis, M.W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. Nat. Neurosci. *21*, 1281–1289. Number: 9 Publisher: Nature Publishing Group, https://doi.org/10.1038/s41593-018-0209-y.
67. Bohnslav, J.P., Wimalasena, N.K., Clausing, K.J., Dai, Y.Y., Yarmolinsky, D.A., Cruz, T., Kashlan, A.D., Chiappe, M.E., Orefice, L.L., Woolf, C.J., et al. (2021). DeepEthogram, a machine learning pipeline for supervised behavior classification from raw pixels. eLife *10*, e63377. https://doi.org/10.7554/eLife.63377.
68. Johnsen, S., Kelber, A., Warrant, E., Sweeney, A.M., Widder, E.A., Lee, R.L., and Hernandez-Andres, J. (2006). Crepuscular and nocturnal illumination and its effects on color perception by the nocturnal hawkmoth Deilephila elpenor. Journal of Experimental Biology *209*, 789–800.
69. Théry, M., Pincebourde, S., and Feer, F. (2008). Dusk light environment optimizes visual perception of conspecifics in a crepuscular horned beetle. Behavioral Ecology *19*, 627–634. https://doi.org/10.1093/beheco/arn024.
70. Striedter, G.F. and Northcutt, R.G. (2020). Brains through time: a natural history of vertebrates (Oxford University Press). OCLC: 1162849672.
71. Peirson, S.N., Brown, L.A., Pothecary, C.A., Benson, L.A., and Fisk, A.S. (2018). Light and

the laboratory mouse. J. Neurosci. Meth *300*, 26–36. Measuring Behaviour 2016, https://doi.org/10.1016/j.jneumeth.2017.04.007.

72. Wang, W., Sun, C., and Chao, H. (1997). Color image segmentation and understanding through connected components. In 1997 IEEE International Conference on Systems, Man, and Cybernetics. Computational Cybernetics and Simulation, volume 2. pp. 1089–1093 vol.2.
73. Chang, C.L. and Lee, R.C.T. (1972). Problem-Solving Methods in Artificial Intelligence (Nils J. Nilsson). SIAM Review *14*, 513–514. https://doi.org/10.1137/1014093.
74. Shannon, C.E. and Weaver, W. (1949). The mathematical theory of communication (Urbana,: University of Illinois Press).
75. Ziebart, B.D., Bagnell, J.A., and Dey, A.K. (2010). Modeling interaction via the principle of maximum causal entropy. Carnegie Mellon University .
76. Mandelbrot, B.B. (1983). The Fractal Geometry of Nature, volume 1 (WH freeman New York).
77. Romero, S., Campbell, J.F., Nechols, J.R., and With, K.A. (2009). Movement behavior in response to landscape structure: the role of functional grain. Landscape Ecology *24*, 39–51. https://doi.org/10.1007/s10980-008-9278-8.
78. Sebok, D., Vasarhelyi, L., Szenti, I., Vajtai, R., Konya, Z., and Kukovecz, A. (2021). Fast and accurate lacunarity calculation for large 3D micro-CT datasets. ACTA MATERIALIA *214*. https://doi.org/10.1016/j.actamat.2021.116970.