



Original Articles

Exploration patterns shape cognitive map learning

Iva K. Brunec^{a,b,*}, Melissa M. Nantais^c, Jennifer E. Sutton^{c,d}, Russell A. Epstein^b, Nora S. Newcombe^a^a Temple University^b University of Pennsylvania^c University of Western Ontario^d Brescia University College

ARTICLE INFO

Keywords:

Spatial navigation
 Cognitive map
 Exploration
 Virtual reality
 Roaming
 Environmental structure
 Space syntax

ABSTRACT

Spontaneous, volitional spatial exploration is crucial for building up a cognitive map of the environment. However, decades of research have primarily measured the fidelity of cognitive maps after discrete, controlled learning episodes. We know little about how cognitive maps are formed during naturalistic free exploration. Here, we investigated whether exploration trajectories predicted cognitive map accuracy, and how these patterns were shaped by environmental structure. In two experiments, participants freely explored a previously unfamiliar virtual environment. We related their exploration trajectories to a measure of how long they spent in areas with high global environmental connectivity (integration, as assessed by space syntax). In both experiments, we found that participants who spent more time on paths that offered opportunities for integration formed more accurate cognitive maps. Interestingly, we found no support for our pre-registered hypothesis that self-reported trait differences in navigation ability would mediate this relationship. Our findings suggest that exploration patterns predict cognitive map accuracy, even for people who self-report low ability, and highlight the importance of considering both environmental structure and individual variability in formal theory- and model-building.

1. Introduction

Spontaneous, volitional spatial exploration is a fundamental feature of human and animal behavior (E. Schulz, Wu, Ruggeri, & Meder, 2019; L. Schulz, 2012; Zurn & Bassett, 2018). Spatial exploration is valuable to an organism because it can facilitate constructing a representation of the spatial structure of the world—a cognitive map—which can then be used to guide goal-directed behaviors such as wayfinding, foraging, and returning to the home base (Calhoun & Hayden, 2015; Davidson & El Hady, 2019; Pfeiffer & Foster, 2013). In the classical formulation, exploration and curiosity are key signatures of a cognitive map. However, the relationship between exploration and cognitive maps is rarely studied, and we know little about how the acquisition of cognitive maps is shaped by different features of exploratory behavior.

Some indirect evidence for a relationship comes from neuroscience. Exploration and active information seeking are prominent drivers of activity in the medial temporal lobes (Gruber & Ranganath, 2019), key

brain regions supporting cognitive maps (Bellmund, Gärdenfors, Moser, & Doeller, 2018; Epstein, Patai, Julian, & Spiers, 2017; Momennejad, 2020; Peer, Brunec, Newcombe, & Epstein, 2021). Spatial exploration involves active movement, either of the body or the sense organs. In rodents, active foraging and exploration produce stronger hippocampal place cell firing compared to passive transportation (Song, Kim, Kim, & Jung, 2005), better coding of spatial information in place cells (Terrazas et al., 2005), and stronger grid-like responses in entorhinal cortex (Winter, Mehlman, Clark, & Taube, 2015). Visual exploration during spontaneous stop-and-scan events produces place cell firing (Monaco, Rao, Roth, & Knierim, 2014), and visual exploration in both primates and humans evokes entorhinal grid-cell-like representations (Julian, Keinath, Frazzetta, & Epstein, 2018; Killian, Jutras, & Buffalo, 2012). Exploration also involves decision making, and human neuroimaging studies have shown that hippocampal activity during naturalistic virtual navigation is stronger when active decision-making is required than when it is absent (Howard et al., 2014; Kaplan et al., 2012; Spiers &

* Corresponding author at: Department of Psychology, University of Pennsylvania, Goddard Laboratories, 3710 Hamilton Walk, Philadelphia, PA 19104, United States of America.

E-mail address: ivab@sas.upenn.edu (I.K. Brunec).

<https://doi.org/10.1016/j.cognition.2022.105360>

Received 11 April 2022; Received in revised form 8 December 2022; Accepted 11 December 2022

Available online 20 December 2022

0010-0277/© 2022 Elsevier B.V. All rights reserved.

Maguire, 2006). Together, these data show that active navigation, visual sampling, and decision making—key elements of exploration—induce the activation of neural structures that mediate cognitive maps. This prompts the question of how exploration supports the formation of such maps (Chrastil & Warren, 2012), and how individual differences in exploratory behavior might explain the strikingly variability of the cognitive maps that different navigators can form of the same environment (Furman, Clements-Stephens, Marchette, & Shelton, 2014; Ishikawa & Montello, 2006; Uttal et al., 2013; Weisberg & Newcombe, 2016; Weisberg, Schinazi, Newcombe, Shipley, & Epstein, 2014; Wolbers & Hegarty, 2010).

In the present study, we investigated the relationship between exploration and cognitive maps. Specifically, we extracted moment-to-moment dynamics of behavior as participants explored a new virtual environment (Virtual Silcton), and then related these exploration patterns to the fidelity of the cognitive maps that the participants formed. Our approach contrasts with that used in the typical cognitive map study, in which spatial memories are assessed by contrasting pre-learning and post-learning ‘snapshots’—a procedure that obscures the rich dynamical process that takes place during real-world learning episodes. We reasoned that if we wanted to understand how exploration allows a navigator to build up a representation of the environment, we needed to unravel the step-by-step process of exploration, the choices people make, where they tend to go, and where they tend to linger.

One aspect of navigational behavior we considered was the extent to which the environment was fully explored. Intuitively, one might consider that an optimal way to explore a novel environment would be to seek out new locations until every possible state in an environment is covered. Variations of this exploration approach have been implemented in machine learning agents (Chentanez, Barto, & Singh, 2004; Vezzani, Gupta, Natale, & Abbeel, 2019), based on the assumption that we are intrinsically motivated to seek out novelty and biased towards visiting previously unseen or less visited spaces. If this is the optimal approach, more accurate spatial memory for a particular environment should stem from higher coverage of the environment. *Roaming entropy* provides a measure of how many different locations, or states, one visits in a given amount of time (Fig. 3A). Maximal roaming entropy would be achieved by spending equal amounts of time in each environmental location. Rodents reared in enriched environments show higher roaming entropy in new environments and higher hippocampal neurogenesis rates compared to rodents reared in impoverished environments (Freund et al., 2013). Further, humans with higher roaming entropy in their day-to-day GPS traces have stronger hippocampal-striatal connectivity relative to those who explore less (Heller et al., 2020). A recent study also found that patients with Alzheimer's Disease show significantly reduced entropy compared to healthy older adults (Ghosh, Puthusserypady, Chan, Mascolo, & Hornberger, 2022). To our knowledge, the inter-individual variability in this measure of exploration has not been linked to the subsequent structure of cognitive maps. A somewhat related measure is *displacement* from the starting point. A previous study showed that participants with lower displacement who revisited the same locations more often had less precise memory relative to those with higher displacement and less revisiting (Gagnon et al., 2018), replicating real world findings (Munio, Stefanucci, Rovira, Squire, & Hendricks, 2019).

Measures of coverage, however, are agnostic to environmental structure, and we know that the geometry of our environments shapes our behavior (Barhorst-Cates, Meneghetti, Zhao, Pazzaglia, & Creem-Regehr, 2021; Coutrot et al., 2022) and neural representations (Krupic, Bauza, Burton, Barry, & O'Keefe, 2015). Thus, a second issue that we considered is how exploration is shaped by the spatial structure of the environment. This structure can be characterized by *space syntax*. This set of spatial analysis methods, which were originally developed in architecture, conceptualize the environment as a graph of connected locations (Hillier & Hanson, 1989; Penn, 2003). Previous work has shown that connectivity of individual street segments predicts people's

gaze patterns and mapping ability (Emo, 2014; Pagkratidou, Galati, & Avraamides, 2020; Walter, Essmann, König, & König, 2022), and that human hippocampal activity when virtually navigating through a city neighborhood scales with step-by-step changes in both local and global connectivity of street segments (Javadi et al., 2017). In the current study, we examined whether the representations that participants form of a new environment are affected by the congruence of their exploration patterns with the spatial syntax of the environment. To do this, we derived a new measure of *experienced integration*, which describes the extent to which participants concentrate exploration on areas of high global connectivity.

We also considered issues related to individual differences in navigational ability. Prior work using the Virtual Silcton environment found that self-reported navigational ability, measured by the Santa Barbara Sense of Direction Scale (SBSOD; Hegarty, 2002), had a moderate but significant positive relationship with subsequent cognitive map accuracy, assessed by asking participants to point between locations (Weisberg et al., 2014). In the present work, we also used absolute pointing error (see Methods) as a measure of cognitive map accuracy. Experiment 1 was exploratory in nature and had no preregistered hypotheses. Based on significant relationships between exploration measures and subsequent cognitive map accuracy in this experiment, we [preregistered hypotheses](#) for Experiment 2, testing whether those who reported better navigational ability in everyday life (as assessed by SBSOD) would differ in their exploratory behavior from those who reported worse navigational ability in everyday life.

To anticipate our results, in Experiment 1, we found that experienced integration significantly predicted participants' ability to point between locations and to draw a map of the environment, but roaming entropy did not predict either measure. Thus, participants who visit more interconnected parts of the environment form better cognitive maps. In Experiment 2, we tested two preregistered hypotheses stemming from Experiment 1. First, we hypothesized a replication of the link between experienced integration and cognitive map accuracy. Second, we predicted that experienced integration would mediate the relationship between trait differences in navigational abilities and cognitive map accuracy. We found evidence for the former, but not the latter, hypothesis. These results suggest that environmental structure shapes participants' cognitive maps, but exploration patterns are not driven by self-reported navigational ability, highlighting the possibility of using directed exploration to construct more accurate cognitive maps.

2. Methods

2.1. Experiment 1

2.1.1. Participants

Data were collected from 84 participants at the University of Western Ontario (31 male). These participants were part of a larger sample of 172 participants, collected for a separate purpose (see Procedure below; Nantais, 2019), and included all participants in the larger sample who had complete navigation logs during exploration. The remaining participants in the larger sample had incomplete navigation logs due to a logging error in the Virtual Silcton software. Participants received course credit or \$15. The study was approved by the University of Western Ontario Non-Medical Research Ethics Board.

2.1.2. Virtual environment

The exploration of Virtual Silcton was implemented in Unity game software (<https://unity.com/>). The environment was presented to participants using the online Virtual Silcton platform (<https://virtuallilcton.com>), which automatically synced data to a server. Participants freely moved around the world by using the arrow keys on the keyboard, and could adjust their view by moving the mouse. In contrast to previous implementations of Virtual Silcton (Weisberg et al., 2014; Weisberg & Newcombe, 2016), participants were not guided along any

predetermined routes, and were able to leave the paths (i.e., there were no invisible walls in the environment).

2.1.3. Procedure

All participants completed three phases of the experiment: 1) exploration, 2) on-site pointing, and 3) final map drawing. As part of the larger study, each participant was assigned to one of four conditions for the exploration phase: Sketch, Silcton Task, Non-Silcton Task, or Baseline. Participants in all groups were given 16 min in total to explore the environment and find all 8 named buildings. In the Sketch condition, participants were asked to pause exploring every four minutes and draw a map of the environment for one minute each time. In the Silcton Task condition, participants were again paused every four minutes and were given a checklist of items which included a mixture of landmarks seen in Silcton and foils. The Non-Silcton condition followed the same pattern, except that participants were given a coloring page for 1 min following each 4-min interval. Finally, in the Baseline condition, participants continuously explored with no breaks.

2.1.4. Cognitive map accuracy

The accuracy of participants' cognitive maps was measured in two ways: 1) absolute pointing error in an on-site pointing task and 2) Gardony scores of the maps drawn by the participants. In the pointing task, participants were placed into the virtual environment in front of each of the buildings. They were asked to point in the direction of every other building by using their cursor to adjust their viewing direction and clicking to record their response. Absolute pointing error was calculated as the average absolute deviation, in degrees, from the correct direction. The Gardony map score was calculated for the participants' hand-drawn maps of the environment using the Gardony Map Drawing Analyzer software (Gardony, Taylor, & Brunyé, 2016). The resulting measure is a correlation ranging from 0 to 1, where 0 indicates no correspondence between the drawing and the ground truth, and 1 indicates perfect correspondence.

The 78 participants included in the present study were relatively evenly distributed across the four learning conditions (20, 21, 20, and 17 in the Sketch, Silcton Task, Non-Silcton Task, and Baseline conditions, respectively). For the main analyses reported below, we collapsed participants across the four learning conditions as these were not relevant to our current hypothesis. Supplementary analyses found no significant effect of learning condition on absolute pointing error ($F < 1$) but a significant relationship between learning condition and Gardony scores ($F(3, 74) = 2.98, p = .037$). However, when we included *learning condition* as a covariate in the linear model reported below, we found that experienced integration remained a significant predictor ($F(1, 72) = 4.56, p = .036$), while learning condition was not significant ($F(3, 72) = 2.15, p = .100$). Learning condition also remained a non-significant predictor of absolute pointing error in the overall model ($F < 1$). There were also no significant differences in experienced integration scores between the four learning conditions ($F(3, 74) = 1.42, p = .243$).

2.1.5. Dataset subsetting

Out of the full sample of 172 participants, 5 were removed for falling >2 standard deviations below the mean in the pointing task, 1 due to experimenter error, and 2 participants for not disclosing their gender. While the present study made no predictions regarding gender differences, we were unable to recover the data from these two participants. Of the 164 participants remaining after these exclusions, 78 had complete navigation logs. The other participants had complete pointing and map-drawing datasets but due to a technical issue with the Virtual Silcton website, their navigation logs were truncated, meaning that each contained no more than a few minutes of data.

2.2. Experiment 2

2.2.1. Participants

We recruited 52 participants to test the preregistered hypotheses (<https://osf.io/8dj5x>). Participants for this sample were recruited on SONA from the Temple Psychology and the University of Pennsylvania undergraduate research pools. All participants completed the task for credit. The experiment was approved by both university Institutional Review Boards.

2.2.2. Procedure

Participants completed the exploration task on the same online platform as in Experiment 1. Due to the Covid-19 pandemic, we were unable to test participants in person. Instead, participants joined a Zoom call with their camera off and shared their screen. This enabled the experimenter to provide participants with verbal instructions, answer any questions in real-time, and monitor progress throughout the task. The participants' screens were not recorded during these sessions, but their locations when navigating in the virtual reality environments were logged.

All participants completed three phases: 1) exploration, 2) on-site pointing, and 3) model-building. During exploration, each participant was given 25 min to roam the environment freely. They were instructed to attempt to find the 8 named buildings and to remember the names and locations of these buildings. They explored the environment continuously, but they were alerted every 5 min to help them keep track of time. The on-site pointing task was the same as in Experiment 1: participants were placed next to each building in the virtual environment and asked to point to every other building. Because the study was completed remotely, we were not able to collect drawn maps. Instead, we asked the participants to complete a drag-and-drop model-building task. Participants were given a blank rectangle on the screen and bird's eye view images of all eight buildings. They were instructed to drag and drop the images on the screen in the correct configuration to construct a map of the environment. The on-site pointing task and the model-building task are described in detail in Weisberg et al. (2014; 2015). Both tasks were untimed.

We also asked participants in this experiment to complete the Perspective-Taking Task (for Adults; Frick, Möhring, & Newcombe, 2014) and the Santa Barbara Sense of Direction Task (Hegarty, 2002). These standardized measures of spatial ability were included to test whether the exploration patterns we observed were related to trait differences in navigational ability. Participants filled out the SBSOD questionnaire prior to participating to avoid biasing their responses based on their perceived accuracy in the present task, and completed the PTT-A task at the very end of the experiment. We did not analyze the PTT-A results in the present experiment.

2.2.3. Cognitive map accuracy

Accuracy in the on-site pointing task was assessed using the same procedure as Experiment 1. For the drag-and-drop model building task, we used bidimensional regression to assess the consistency between the coordinates of the placed icons and the locations of the corresponding buildings in the environment. The resulting R^2 value (ranging from 0 to 1) was used as the dependent variable.

2.3. Data preprocessing

Prior to running any analyses, we removed all periods during which participants were stationary for >30 s from the navigation logs. This approach served to filter out the 1-min time periods when participants were completing tasks associated with their assigned condition in between periods of exploration in Experiment 1. Pointing scores collected via virtualsilcton.com in Experiment 1 were corrected using the algorithm described in (Weisberg, Schinazi, Ferrario, & Newcombe, 2022). Pointing scores in Experiment 2 did not need to be corrected.

2.4. Data analysis

To test the relationship between experienced integration and subsequent cognitive map accuracy, we performed linear regression analyses in R (R Core Team, 2018). We report standardized beta estimates extracted using the `lm.beta` package in R (Behrendt, 2014). Effect sizes are expressed as partial eta squared values (η_p^2) for each predictor (calculated using the `effectsize` package; Ben-Shachar, Lüdtke, & Makowski, 2020).

2.5. Experienced integration

To relate environmental characteristics to participants' exploration patterns, we relied on axial integration values previously reported for Virtual Silcton by Pagkratidou et al. (2020). As a summary measure for each participant, we calculated the average integration value of their trajectory (= *experienced integration*). The integration value at each 100 ms time increment (data sampling rate) was added to a vector and averaged across the entire time each participant spent exploring (after filtering out stationary periods of over 30s, as described above). The *experienced integration* value is therefore a single number; the average of moment-by-moment values for each participant (Fig. 1A).

Participants who spent more time on high-integration segments would therefore have higher experienced integration values than those who spent more time exploring low-integration segments. When deriving this measure, all coordinates >10 virtual units away from the path were removed, as no integration value could be assigned to them. As a control covariate, we also calculated the proportion of off-path timepoints out of the total exploration time for each participant.

2.6. Roaming entropy

Roaming entropy was calculated as Shannon's entropy, across 2500 states in the environment (50×50 grid). Roaming entropy is based on the probability of a subject being in a particular state, or location, in a given environment across the entire time period of observation (Freund et al., 2013; Heller et al., 2020).

$$REntropy = \sum_{j=1}^n (p_{ij} \times (\log_2 p_{ij})) / \log(n)$$

In this equation, p_{ij} represents the proportion of time across the entire exploration phase that participant i spent in location j . The total number of states in the environment (in our case, 2500), is represented by n .

2.7. Speed: Mean squared displacement

As an additional measure, we also calculated speed of movement. Speed was expressed as the mean squared displacement (MSD), calculated as the squared value of the displacement at each time-step in any direction, following prior approaches (Gagnon et al., 2018). The step-wise values were then averaged to derive a single value for the entire exploration phase for each participant.

$$MSD = \frac{\sum (j_i - j_{i-1})^2}{n}$$

The speed of movement was fixed (i.e., no acceleration/deceleration), and MSD was calculated in 100 ms timesteps, the rate at which data was logged throughout the experiment. In this sense, MSD provides a measure of pausing during exploration.

3. Results

We first investigated whether the correspondence between exploration patterns and the graph structure of the environment (*experienced integration*) predicted cognitive map accuracy. We next compared

participants' performance to three alternative simulated model navigators. Finally, we related each participant's coverage of the environment (*roaming entropy*) to their cognitive map accuracy.

3.1. Experienced integration

Our hypothesis was that participants who explored areas of high global connectivity would form better cognitive maps. In Experiment 1, higher experienced integration was related to lower absolute pointing error ($\beta = -0.477, p < .001, \eta_p^2 = 0.234$; Fig. 1C) and more accurate hand-drawn maps ($\beta = 0.295, p = .010, \eta_p^2 = 0.09$; Fig. 1C). The proportion of time spent off-path had no effect on pointing error ($\beta = 0.042, p = .680, \eta_p^2 = 0.002$) or Gardony scores ($\beta = -0.044, p = .696, \eta_p^2 = 0.002$).

In Experiment 2, higher experienced integration again predicted lower absolute pointing error ($\beta = -0.347, p = .015, \eta_p^2 = 0.100$), in line with our pre-registered predictions (Fig. 1D). There was again no significant relationship between pointing accuracy and the time spent off-path ($\beta = 0.164, p = .241, \eta_p^2 = 0.030$). Somewhat surprisingly, however, we did not find a significant relationship between experienced integration and the bidimensional regression R^2 value for the drag-and-drop model building task ($\beta = 0.077, p = .589, \eta_p^2 < 0.001$; Fig. 1D), and we observed a somewhat weak, but significant relationship between the proportion of time spent off-path and the R^2 value ($\beta = -0.292, p = .044, \eta_p^2 = 0.080$).

3.2. Simulated navigation trajectories

The results above show that there is a relationship between exploration patterns and cognitive map formation. One possible driver of this relationship might be strategy: participants who are able to integrate across distal locations in the environment (Fig. 1C-D) might recognize the areas in the environment that are more informative and approach these areas with greater frequency compared to peripheral areas. Such a difference in strategy might be reflected in a tendency to choose paths at decision points that gravitate towards high-connectivity areas. In contrast, less accurate mappers might make choices at decision points that are more random. To formally model this idea, we generated simulated trajectories based on three different possible approaches to exploration: 1) random walk, 2) proportional connectivity bias, and 3) consistent connectivity bias. The code to generate these simulations is available on GitHub (https://github.com/ivabrunec/nav_free_exploration).¹

The proportional and consistent connectivity bias simulations assume that the navigator can see alternative options at each intersection, or that the alternatives can be retrieved from memories formed during prior passage through the same area. However, we did not formally model the contributions of memory and perception in the present approach. The proportional connectivity approach simply assumes a tendency to prefer more connected nodes, while the consistent connectivity approach assumes a tendency to consistently take paths that lead to more connected nodes.

3.2.1. Random walk

It is important to consider what random behavior is for a realistic human navigator. The simplest possibility would be to compare participants' trajectories against Brownian motion patterns, in which they randomly reorient at each step. However, it is unlikely that human navigators, even very bad ones, would exhibit such behavior in a complex, structured environment with clearly defined paths. Instead, to generate 'human-like random behavior', we generated simulated

¹ Note that the precise values will change slightly if the simulation code is re-run, since the generated integration values will differ, as will the sampling in the bootstrapping procedure.

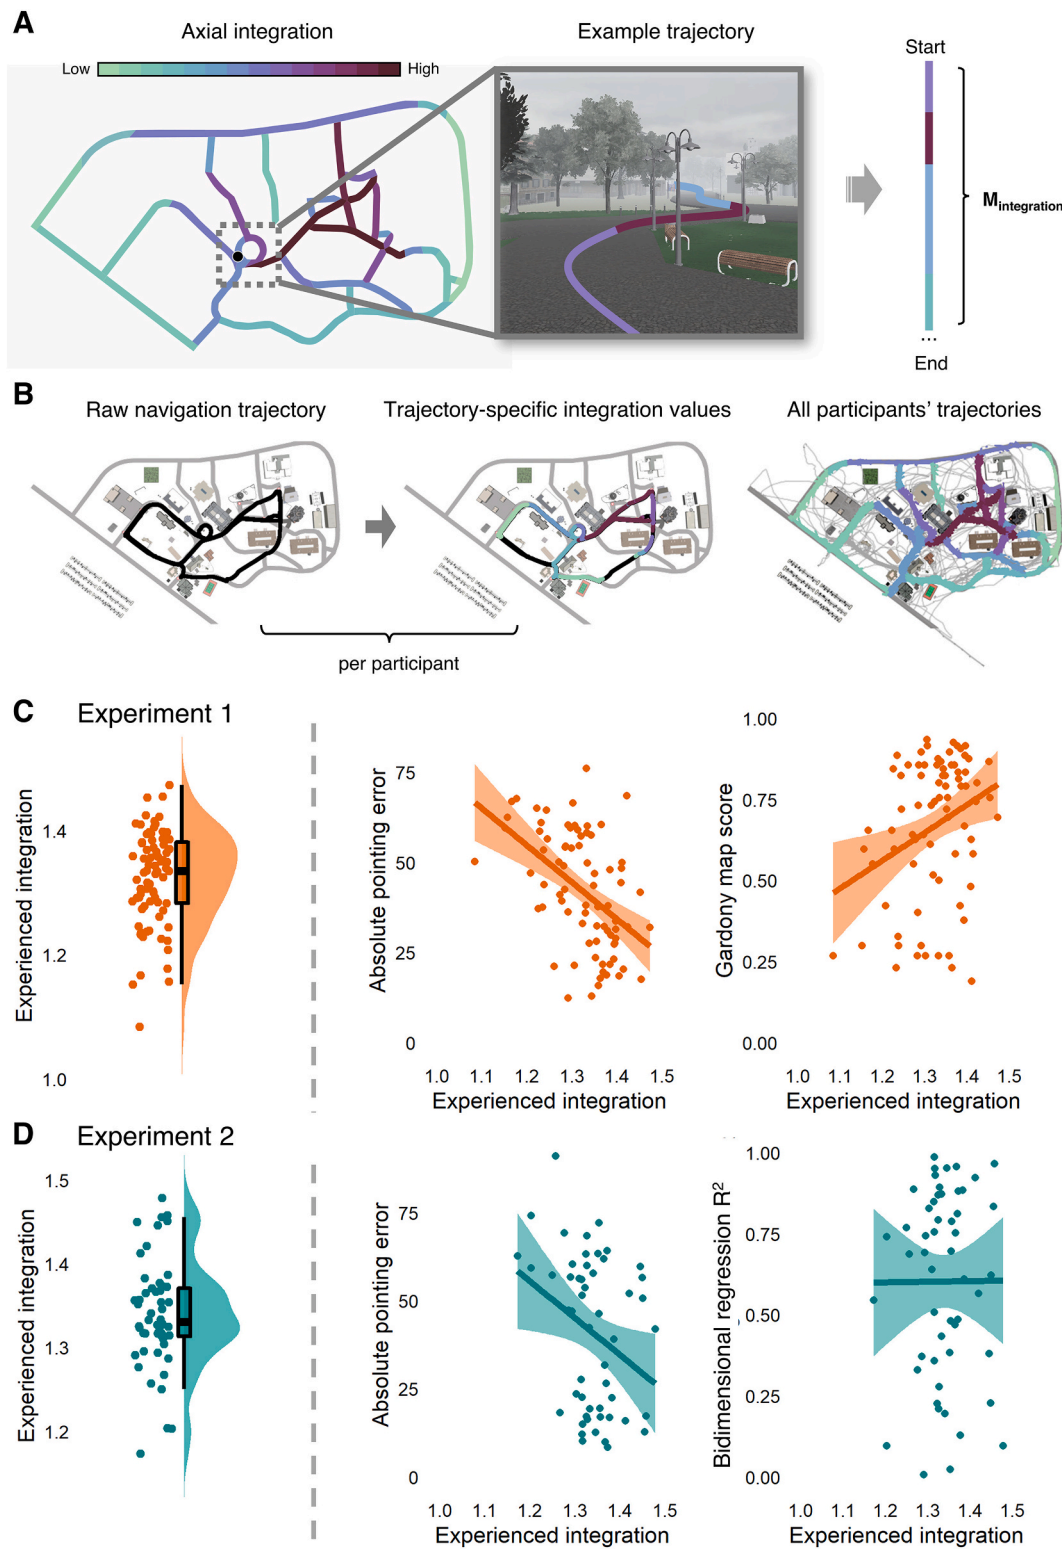


Fig. 1. A) To calculate experienced integration for each participant, we extracted the axial integration values reported by Pagkratidou et al. (2020) for each point along the participant's trajectory and computed their average. The black dot indicates the start location for all participants. B) An example participant's trajectory, color-coded according to the axial integration values, alongside all participants' color-coded trajectories. C) Variability of experienced integration in Experiment 1 (left) was significantly related to participants' pointing ability (middle) and to the accuracy of the maps they drew (right). D) Variability of experienced integration in Experiment 2 (left) was significantly related to participants' pointing ability (middle), but was unrelated to the accuracy of their drag-and-drop maps (right). Panel B was visualized using matplotlib (Hunter, 2007), and the distribution plots in panels C and D were visualized using raincloud plots in R (Allen et al., 2021). The raincloud plots represent the full distribution of the data in scatterplots, and the accompanying boxplots use default ggplot2 settings: the middle notch represents the median, with the upper and lower bounds of the box representing the 25th and 75th quintiles, and the whiskers extending to 95% confidence intervals.

trajectories following random walks through the environment expressed as a graph. Random walks were implemented by making a random choice at each decision point, with the constraint that backtracking was prohibited to prevent the navigator from bouncing between pairs of graph nodes.

3.2.2. Proportional connectivity bias

The second approach we implemented was more directed and comprised weighing the options at each decision point based on their subsequent interconnectivity. Under this strategy, a navigator uses memories for local navigational information (e.g. which nodes have more paths) to make choices that will hopefully guide them to areas of higher global connectivity, and thus allow them to build up a better cognitive map. If we take an example of a decision point with two options, *path 1* leading to a decision point with three connections would be given a weight of 3, while *path 2* leading to a decision point with two connections would be given a weight of 2. Under a proportional connectivity bias, the simulated navigator picks *path 1* 60% of the time and *path 2* 40% of the time (rather than 50/50).

3.2.3. Consistent connectivity bias

In our third approach, the simulated navigator was highly biased to choose the path leading to the most interconnected decision point. Weights were determined the same way as in the approach above, but in this case the navigator consistently chose the path leading to the most interconnected node. If there was more than one option with the highest weight, the navigator randomly selected between the maximally-weighted options. The choices made in both the proportional connectivity bias and the consistent connectivity bias simulations are independent of the navigator's history, but both approaches assume that the navigator has some knowledge or intuition about the environment.

For each of these simulation approaches, we generated 1000 trajectories for each experiment, which matched the average human trajectories in travel time and distance (see Supplementary Materials). We then calculated the average experienced integration for each of the simulated trajectories in the same way as we did for human participants. To relate the exploration behavior of the simulations to the exploration behavior of the human participants, we split participants into tertiles depending on their performance on the onsite pointing task (Fig. 2A-B). This split was not motivated by any belief that participants belong to different clusters, but rather was done to facilitate statistical comparisons between participants at different levels of cognitive mapping ability and the simulations. We then carried out 1000 bootstrapped independent-samples *t*-tests with 1000 samples each (Canty & Ripley, 2021; Davison & Hinkley, 1997) comparing experienced integration for participants in each of the tertiles against experienced integration for a matched number of randomly sampled simulated trajectories, and calculated the proportion of significant *t*-tests in each case. A non-significant *t*-test would indicate that human trajectories in a tertile could not be reliably differentiated from a given simulation approach.

3.3. Experiment 1: Comparison of model to human data

In Experiment 1, human participants had higher integration values than the *random walk* simulation in 98.7% of comparisons for the top tertile, 86.5% of comparisons for the middle tertile, but only 58.8% of comparisons for the bottom tertile (Fig. 2A). These data suggest that while the majority of human participants outperform a random walker, there is a subset of participants (i.e. those in the bottom tertile of mapping performance) whose exploration behavior matches the profile of making a random decision at each decision point.

An even more striking pattern was observed when we compared human participants to the *proportional connectivity bias* simulations. Human participants had higher integration values than this simulated navigator in 73.5% comparisons for the top tertile, but only 34.7% of comparisons for the middle tertile and 7.7% of comparisons for the

bottom tertile.

Finally, we found that human participants had higher integration values than the *consistent connectivity bias* simulation in only 23.3% of the top tertile comparisons, 4.0% of the middle tertile comparisons, and 0.2% of the bottom tertile comparisons. These findings suggest that most human participants, even those with good mapping ability, fail to match the profile of a highly biased navigator who consistently detects and selects the most highly interconnected nodes in the underlying graph of the environment. This result is not surprising, as this latter simulation approach is based on the navigator possessing full knowledge of the environmental structure throughout the exploration phase.

3.4. Experiment 2: Comparison of model to human data

A similar pattern overall was observed in Experiment 2. Human integration values were significantly higher than the *random walk* integration values in 96.2% of comparisons for the top tertile, 96.1% of comparisons for the middle tertile, but only 70.6% of comparisons for the bottom tertile (Fig. 2).

Relating human values to the *proportional connectivity bias* simulations revealed that human navigators had higher integration values than the simulated navigators in 70.2% of comparisons for the top tertile, 66.8% for the middle tertile, and 21.4% for the bottom tertile (Fig. 2A). Thus integration values for participants with middling-to-poor navigational abilities are consistent with an exploration strategy of gravitating towards more interconnected nodes in the underlying graph, but not doing so consistently.

Finally, compared to the *consistent connectivity bias*, higher integration values were observed in 19.9% of comparisons with the top tertile, 18.4% for the middle tertile, and 1.8% for the bottom tertile in Experiment 2 (Fig. 2C).

3.5. Experiment 1 & 2: Comparison to model sample means

To generate group-level statistics, we took the following approach: we sampled each of the simulated datasets 1000 times, randomly selecting as many simulated values as participants in each study in each sample (78 in Experiment 1 and 52 in Experiment 2). We then calculated the mean of each sample and plotted the distribution against the human mean in each tertile. In this way, the proportion of samples that are below the true mean for each human participant tertile can be interpreted as a *p*-value (Fig. 2C).

In Experiment 1, the mean integration values in all three tertiles were significantly greater than the random walk simulation distribution (all *p*-values < .001; Fig. 2C). Compared to the proportional connectivity bias simulation, the means of the participants in the top tertile ($p < .001$) and middle tertile ($p = .005$) were significantly greater than the simulations, while the bottom tertile mean did not significantly differ from this simulation ($p = .545$). The top tertile mean was significantly greater than the consistent connectivity bias as well ($p < .001$), while the middle tertile ($p = .829$) and the bottom tertile ($p > .99$) means were not.

Note that this analysis simply compared the means of the distributions and did not take into account the spread of values within each tertile and each sample of the simulation. This approach is different from our earlier analysis, in which we performed individual statistical tests on sampled distributions. Thus, it is possible to observe a significant difference in means in the current analysis even though in the previous analysis there were only a relatively small number of individual samples that reached significance. For example, the mean of human performance for the top tertile was greater than the mean of the consistent connectivity bias simulation, even though this difference only reached significance in 23.3% of individual samples.

In Experiment 2, the top two tertiles had nearly identical means (Fig. 2C). This was reflected in the significance values: all three tertile means were significantly greater than the random walk simulation (all $p < 0.001$). In contrast to Experiment 1, all three tertile means were

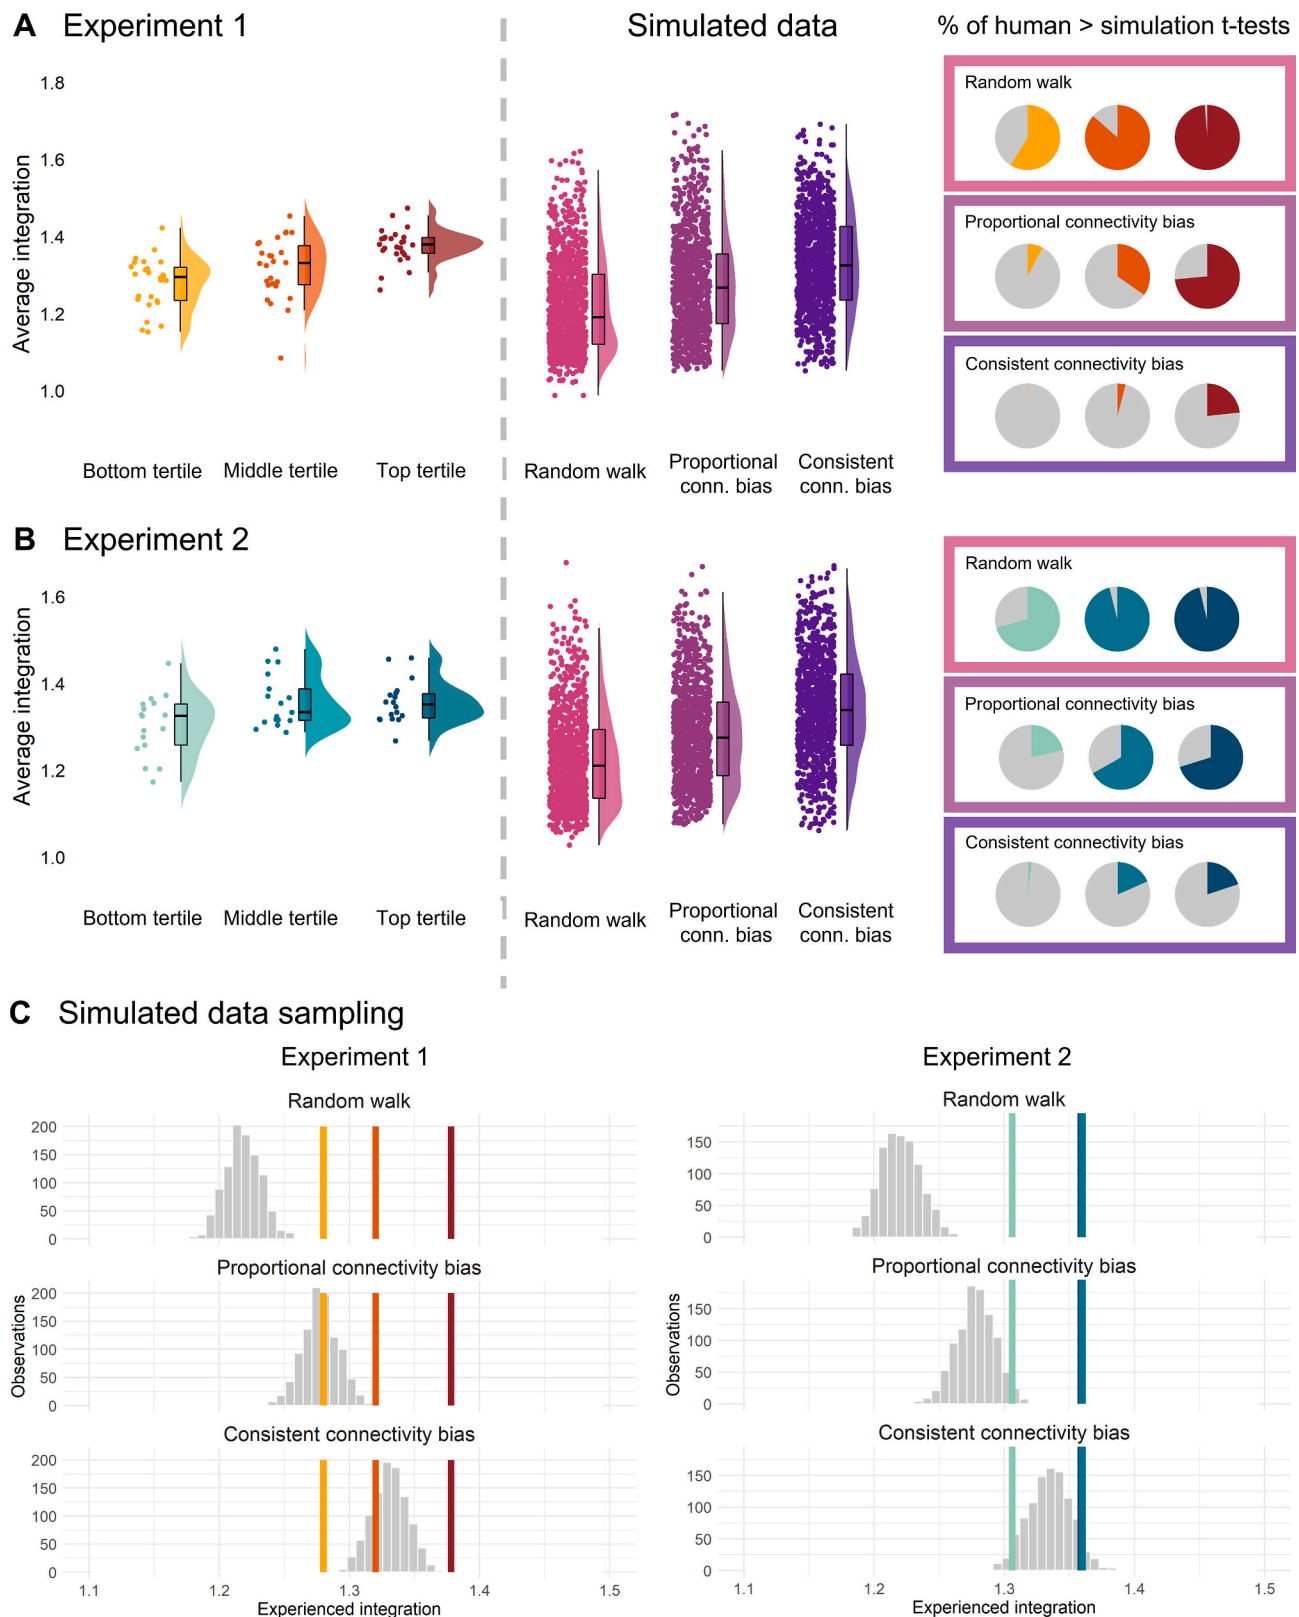


Fig. 2. A) The integration values for each participant in Experiment 1 split into tertiles according to on-site pointing ability (left); integration values for simulated navigators using different strategies (middle); proportion of 1000 bootstrapped *t*-tests for which the human navigators in each tertile perform better than the simulated navigators. B) The same data for Experiment 2. C) The distribution of means sampled from each simulated distribution (1000 samples per mean) vs. the mean experienced integration in each human sample. The same colors denote the same tertiles, always following left-to-right from the bottom to the top tertile. Note that in Experiment 2, the middle and top tertile means were nearly identical, making them difficult to differentiate.

also significantly greater than the proportional connectivity bias simulation, though this effect was more pronounced for the top two tertiles ($p < .001$), compared to the bottom tertile ($p = .042$). As in Experiment 1, the bottom tertile mean was not significantly different from the consistent connectivity bias simulation ($p = .978$), but the middle ($p = .085$) and top tertiles ($p = .095$) were also not significantly different, though consistent with the trend in Experiment 1 (Fig. 2C).

Together, these data make it plausible that the individual variability that we observed in cognitive maps might reflect the use of different navigational strategies that vary in how successfully they draw navigators to regions of high integration.

3.6. Relationship to self-reported navigational ability

How do the individual differences in exploration that we observe relate to trait-level individual differences? Individual differences in navigational ability, measured by the Santa Barbara Sense of Direction scale (Hegarty, 2002) have previously been related to behavioral measures of navigation testing the ability to use cognitive maps. This raises the possibility that they might also relate to the ability to learn new cognitive maps through exploration.

Based on the results of Experiment 1, we preregistered the hypotheses that 1) self-reported navigational ability (SBSOD) should predict pointing ability, and 2) experienced integration should be a mediator of the relationship between SBSOD and pointing ability. To address hypothesis (1), we fit two multiple regression models, predicting absolute pointing error or the model-building R^2 value from experienced integration, SBSOD score, and the time spent off-path.

Experienced integration was a significant predictor of absolute pointing error in the multiple regression ($\beta = -0.358, p = .010, \eta_p^2 = 0.104$), recapitulating our previous finding from the correlation analysis. SBSOD was also a significant predictor, but had an approximately 50% smaller effect size ($\beta = -0.272, p = .049, \eta_p^2 = 0.051$). The time spent off-path was not a significant predictor ($\beta = 0.240, p = .091, \eta_p^2 = 0.058$). Applying the same model to the bidimensional regression R^2 values, we found no significant effect of experienced integration ($\beta = 0.090, p = .502, \eta_p^2 < 0.001$), but there was a significant effect of SBSOD ($\beta = 0.350, p = .012, \eta_p^2 = 0.071$), as well as a significant negative effect of the proportion of time spent off-path ($\beta = -0.390, p = .007, \eta_p^2 = 0.141$). These data suggest that both experienced integration and SBSOD predicted cognitive map accuracy independently, insofar as cognitive map accuracy is indexed by pointing accuracy.

We next performed a mediation analysis predicting pointing accuracy from SBSOD, with experienced integration as a mediator. The analysis was implemented using the *mediation* package in R (Tingley, Yamamoto, Hirose, Keele, & Imai, 2017), running 1000 simulations with nonparametric bootstrapping. We found no significant evidence for a mediation effect ($\beta = -0.191, 95\% \text{ CI } [-2.95, 2.24], p = .856$). We also found no significant evidence for a similar mediation effect on the bidimensional regression R^2 from model-building ($\beta = \sim 0, 95\% \text{ CI } [-0.001, 0.02], p = .802$).

Together, these analyses suggest that experienced integration remains predictive of pointing accuracy when accounting for self-reported trait individual differences in navigational ability, but does not mediate the relationship between trait abilities and cognitive map accuracy. This suggests that exploration efficiency may not be captured by self-reported measures of navigational abilities.

3.7. Roaming entropy

Roaming entropy measures the extent to which a navigator covers the entire environment in their exploration. Despite starting at the same location and being given the same amount of time to navigate, participants showed a range of roaming entropy values. Some participants tended to remain in fewer locations in the environment for a longer period of time, while others covered a greater range in the same amount

of time (Fig. 3A).

Roaming entropy was not a significant predictor of pointing error in Experiment 1 ($\beta = -0.002, p = .986, \eta_p^2 < 0.001$) or Experiment 2 ($\beta = 0.216, p = .125, \eta_p^2 = 0.046$). Nor did it predict Gardony scores in Experiment 1 ($\beta = -0.070, p = .542, \eta_p^2 = 0.005$; Fig. 3B) or bidimensional R^2 values in Experiment 2 ($\beta = -0.215, p = .126, \eta_p^2 = 0.046$; Fig. 3C). These data suggest that the overall coverage of the environment is not, by itself, predictive of better cognitive maps. Rather, it is greater coverage of segments which are more integrated that predicts cognitive map accuracy.

To further understand the effects of experienced integration and roaming entropy, we entered both exploration measures into a single regression model, along with the proportion of time spent off-path. In Experiment 1, pointing accuracy was significantly predicted by experienced integration ($\beta = -0.520, p < .001, \eta_p^2 = 0.240$), but not roaming entropy ($\beta = -0.173, p = .113, \eta_p^2 = 0.028$). The proportion of time spent off-path was also not a significant covariate ($\beta = 0.083, p = .429, \eta_p^2 = 0.008$). Gardony scores were again significantly predicted by experienced integration ($\beta = 0.303, p = .011, \eta_p^2 = 0.091$), but not roaming entropy ($\beta = 0.030, p = .802, \eta_p^2 < 0.001$) or the proportion of time spent off-path ($\beta = -0.051, p = .662, \eta_p^2 = 0.003$).

In Experiment 2, experienced integration had a somewhat smaller effect size in the overall model to Experiment 1, and was a marginally significant predictor of pointing accuracy ($\beta = -0.332, p = .052, \eta_p^2 = 0.097$). Roaming entropy was again not a significant predictor of pointing accuracy ($\beta = 0.031, p = .868, \eta_p^2 = 0.015$), nor was the proportion of time spent off-path ($\beta = 0.146, p = .412, \eta_p^2 = 0.015$). Bidimensional R^2 values were not significantly predicted by any of the measures: experienced integration ($\beta = 0.030, p = .861, \eta_p^2 < 0.001$), roaming entropy ($\beta = -0.094, p = .616, \eta_p^2 = 0.054$), or time off-path ($\beta = -0.236, p = .169, \eta_p^2 = 0.035$).

Finally, we wanted to account for the effect of movement speed on these measures. We added movement speed, calculated as the mean displacement over time, to the overall regression models reported above. In Experiment 1, pointing accuracy was significantly predicted by both experienced integration ($\beta = -0.393, p < .001, \eta_p^2 = 0.260$) and movement speed ($\beta = -0.304, p = .005, \eta_p^2 = 0.102$) when all variables were entered into a single model. As before, roaming entropy and the amount of time spent off-path were not significant (both p -values $> .4$). Gardony map scores were significantly predicted only by experienced integration ($\beta = 0.271, p = .036, \eta_p^2 = 0.092$; all other p -values $> .5$). In Experiment 2, only displacement speed was a significant predictor of pointing accuracy when all of the factors were entered into a single model, suggesting that movement speed may interact with other exploration measures ($\beta = -0.321, p = .030, \eta_p^2 = 0.105$). No other measure was significant (all p -values $> .3$). No measure was significantly predictive of the model-building R^2 coefficient (all p -values $> .2$).

4. Discussion

In the present study, we investigated how humans explore a previously unfamiliar virtual reality environment and how these exploration patterns affect the accuracy of the cognitive maps that they form. We found that experienced integration, or the amount of time spent in highly-interconnected parts of the environment, predicted cognitive map accuracy in two experiments, whereas roaming entropy, or the overall coverage of the environment, did not.

While self-reported navigational ability was weakly related to accurate pointing ability, it did not relate to experienced integration, and experienced integration did not mediate the relationship between self-reported navigational ability and spatial memory. It is also worth noting that the relationship between individual differences in experienced integration and cognitive mapping was somewhat weaker in Experiment 2. Together, these data illustrate how different exploration patterns give rise to variability in cognitive maps and suggest that the structure of the environment plays an important role in the structure of

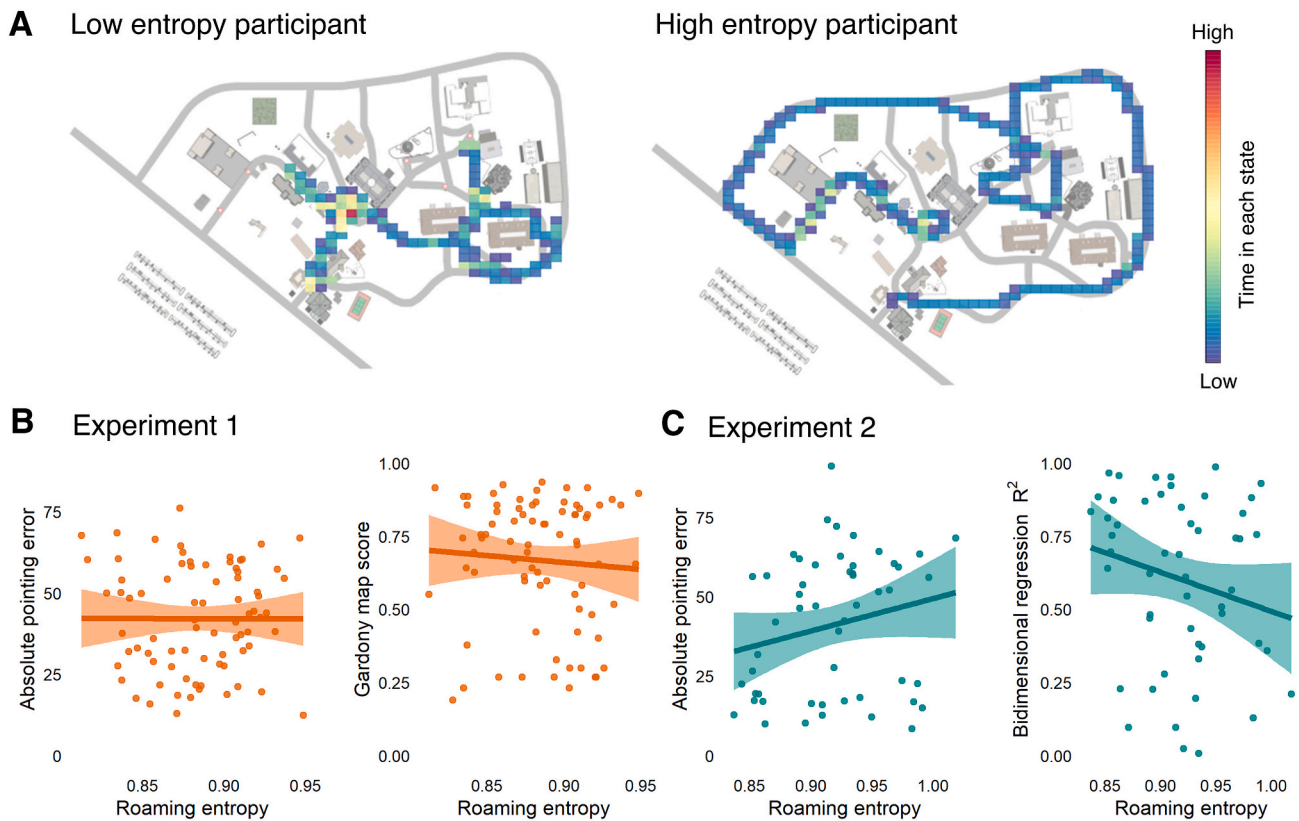


Fig. 3. A) Example low and high entropy participants. Roaming entropy was not a significant predictor of pointing accuracy, hand-drawn map accuracy, or model-building accuracy in B) Experiment 1 or C) Experiment 2.

acquired knowledge.

In the present study, participants were given no specific guidance on how to explore the environment, they were simply instructed to find all 8 buildings and remember their names. Despite the open-ended nature of this task, we found consistent differences in strategies between those participants who were able to build a relatively accurate map of the environment, and those who were not. Put simply, better mappers' exploration patterns tended to track the network structure of the environment more closely than the patterns of less accurate mappers (*experienced integration*), and they tended to pause less while navigating (*mean squared displacement*). These findings suggest that moment-to-moment behavioral patterns can provide a window into the likely nature of the subsequent spatial representation.

Our findings replicate prior work suggesting that navigators with higher levels of spatial exploratory behavior form more accurate cognitive maps (Farran, Blades, Hudson, Sockeel, & Courbois, 2022; Gagnon et al., 2018; Munion et al., 2019). Also mirroring prior results, we found no relationship between exploration measures and self-reported navigational abilities (Gagnon et al., 2018), despite significant relationships between 1) exploration and pointing performance and 2) self-reported spatial abilities and pointing performance. The present study also replicates the observation that participants who navigated more consistently and paused less and for shorter periods of time subsequently showed higher spatial memory performance (Gagnon et al., 2018).

Higher levels of exploration have also been linked to lower spatial anxiety (Gagnon et al., 2018; Gagnon, Cashdan, Stefanucci, & Creem-Regehr, 2016) and higher well-being in general, as measured by day-to-day experience sampling (Heller et al., 2020). These findings converge on questions to be addressed in future research: are individuals with higher exploration tendencies less anxious in general, and does the tendency to explore in the spatial domain translate to curiosity and

information sampling in non-spatial domains? This notion would fit with a previous finding that trait-level curiosity is predictive of knowledge network-building patterns in participants exploring Wikipedia pages (Lydon-Staley, Zhou, Blevins, Zurn, & Bassett, 2021). Further, recent work suggests that both spatial and conceptual learning are captured by the same Bayesian generalization model (Wu, Schulz, Garvert, Meder, & Schuck, 2020), but it is unclear whether this would also extend to spatial knowledge acquired by active first-person navigation.

4.1. Simulated trajectories

Our simulation approach suggested that the navigation patterns of participants with different cognitive mapping abilities can be captured by the use of different strategies at decision points. Even the least accurate mappers' exploration patterns did not tend to be random, but rather consistent with a slight preference for interconnected nodes. In contrast, the most accurate mappers' patterns were most akin to a directed walk, preferring to take paths that lead to nodes with highest connectivity values. These findings suggest that humans tend to select paths with many interconnections, relying on either visual information or memory for previously visited locations. The balance between these alternatives remains to be explored in future studies, but recent work suggests that even eye movement patterns during navigation are reflective of the graph structure of the environment (Walter et al., 2022).

Importantly, the similarity between human and simulated trajectories suggests that participants were relying on a *forward-looking* search strategy, as the decision of which path segment to take was reached based on the connectivity value of the *subsequent* decision point. These data are broadly consistent with the notion that mammals form predictive maps by extracting regularities and structure from the environment (De Cothi et al., 2022; Momennejad, 2020; Stachenfeld,

Botvinick, & Gershman, 2017). Beyond the convergence with this theoretical viewpoint, our results also highlight the importance of carefully characterizing human navigational behavior when building computational models or artificial agents (Devlin et al., 2021) and for better understanding human exploration tendencies in new environments.

The structure of the task at hand and the structure of the environment are important drivers of learning. The nature of connections between bits of information shapes the speed of learning and nature of representations humans can form (Karuza, Thompson-Schill, & Bassett, 2016; Lynn, Kahn, Nyema, & Bassett, 2020; Qian et al., 2022; Solway et al., 2014). For example, when information is organized as a structured graph, humans tend to pick up on clusters that are closest together on relevant dimension, such as time, space, or semantic distance. Establishing these clusters, or communities, enables us to represent tasks hierarchically and flexibly (Karuza et al., 2016; Solway et al., 2014) as we build structured knowledge by accumulating pieces of information over time (Duncan & Schlichting, 2018; Schapiro, Rogers, Cordova, Turk-Browne, & Botvinick, 2013; Schapiro, Turk-Browne, Norman, & Botvinick, 2016). Simulations and human behavioral experiments suggest that the structure of such information graphs is most useful to us if the communities are tightly clustered into ‘neighborhoods’ of information, highlighting that environments or knowledge structures with different levels of complexity also have different levels of learnability (Karuza et al., 2016; Lynn et al., 2020).

4.2. Correspondence between environmental features and exploration measures

In the spatial domain, the geometry of the environment that one grew up in appears to affect subsequent navigation and spatial learning abilities (Barhorst-Cates et al., 2021; Coutrot et al., 2022). However, it has not yet been systematically tested whether individuals are similarly affected on all new spatial learning or whether there are gradations in strategy *flexibility*, rather than general *ability*. Future work will be necessary to determine whether individuals who are successful in a given exploration task would succeed at wayfinding in any new environment, or whether there are systematic interactions between individuals’ ability and environmental structures. A recent investigation of millions of participants suggests that growing up in an environment with a more regular structure impairs the ability to flexibly integrate across locations in new complex environments (Coutrot et al., 2022). Evidence from non-spatial associative structure learning in humans suggests that different stimulus presentation sequences corresponding to different walks through the graph of the task shape the observed behavioral biases (Karuza, Kahn, Thompson-Schill, & Bassett, 2017). Moreover, when participants are able to freely navigate such non-spatial graphs, their planning reflects the underlying hierarchical structure of the task (Solway et al., 2014; see also Balaguer, Spiers, Hassabis, & Summerfield, 2016). These observations dovetail with our present findings which suggest that participants generally showed sensitivity to the connectivity between the nodes in the present environment. However, we also present evidence that there are substantial individual differences in how efficient participants are at extracting such structure, and formal models of learning and exploration should account for these differences.

Interestingly, in the present experiment, the overall coverage of the environment (roaming entropy) was not predictive of cognitive map accuracy, again suggesting that the connectivity of the world was more relevant than the total amount of new territory covered. However, it is important to consider that the environment used in the present study was highly structured, with clearly defined paths. In such environments, the underlying graph properties may be the key factor constraining navigational strategy (Wiener & Mallot, 2003; Wiener, Schnee, & Mallot, 2004). In large-scale open environments, on the other hand, navigational strategies fundamentally shift due to the difference in reference frames (Wolbers & Wiener, 2014), and direction-free displacement

measures, such as diffusion or roaming entropy, might capture more navigational variance. Future studies should compare how well different measures capture exploration success in different environmental structures, and examine whether better navigators are able to adapt their strategy flexibly across these environments.

4.3. Limitations of the present work

The present study had two key limitations. First, the effects we observed in Experiment 2 were somewhat attenuated relative to Experiment 1. In contrast to Experiment 1, we observed no significant relationship between exploration measures and the R^2 value in the model-building task. There are multiple possible reasons for this discrepancy, but it is worth noting that participants were explicitly told that they would have to draw a map in Experiment 1 but not in Experiment 2. The absence of such a requirement may have led Experiment 2 participants to focus on encoding the buildings and their names, but not the relationships between them. Participants were also given more time to explore in Experiment 2 (25 min vs. 16 min) and each person completed the task on their own personal computer. While the participants were on a video call with the experimenter throughout the entire experiment, remote data collection meant that we were unable to standardize aspects such as screen size and resolution, distance from the screen, and potential brief glitches due to network/processor issues. Anecdotally, the longer exploration time in Experiment 2 meant that some participants indicated that they were ‘done’ exploring before the time was up, and were prompted by the experimenter to continue. This may have contributed to more aimless roaming in this experiment and diluted the strength of the effect of initial exploration.

Second, while we observed an otherwise consistent effect of environmental structure on cognitive map accuracy, we did not manipulate it, nor did we direct participants along more vs. less integrated paths. Because we did not manipulate the structure of the environment, we cannot be confident that the differences in strategies between more and less accurate mappers are not incidental. Nonetheless, the nature of the observed exploration patterns shaped subsequent performance, even if the exploration patterns themselves were incidental. Future studies will be necessary to disentangle the contributions of *active* navigational decision-making vs. the optimality of the route one experiences in a given environment. For example, it might be the case that simply being guided along inter-connected paths is sufficient for forming an accurate cognitive map. Alternatively, it is possible that less accurate mappers are also less efficient at extracting environmental structure from ongoing experience. This issue is important because it relates to the potential efficacy of guidance and signage designed to aid wayfinding in the real world.

4.4. Conclusions

In everyday life, we often actively seek out new information and connections between existing elements of knowledge through exploration – be it in space, social contexts, or existing elements of semantic knowledge (Lydon-Staley et al., 2021). According to some estimates, we consume approximately 34 gigabytes of information in media alone on a daily basis (Bohn & Short, 2012), which is a daunting amount of content to explore, but somehow we are able to seamlessly organize and structure the new knowledge acquired through everyday experience. The present findings suggest that, in the spatial domain at least, this challenge is accomplished in part by the use of intelligent information sampling strategies, and that fine-grained differences in these sampling strategies are reflected in subsequent cognitive maps.

CRedit authorship contribution statement

Iva K. Brunec: Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing – original draft, Visualization. **Melissa**

M. Nantais: Conceptualization, Methodology, Formal analysis, Investigation, Resources, Writing – review & editing. **Jennifer E. Sutton:** Conceptualization, Methodology, Resources, Writing – review & editing, Supervision, Project administration, Funding acquisition. **Russell A. Epstein:** Conceptualization, Methodology, Writing – review & editing, Supervision, Project administration, Funding acquisition. **Nora S. Newcombe:** Conceptualization, Methodology, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Data availability

The link to the data and code is provided in the paper.

Acknowledgments

This research was supported by a Discovery grant from the Natural Sciences and Engineering Research Council to J.E.S., National Institutes of Health grant R01 EY031286 to R.A.E., and National Science Foundation grant EHR 1660996 to N.S.N. We thank Sarah Hendricks and Mitchell Decasere for their help with data collection for Experiment 2. We thank Marianna Pagkratidou, Alexia Galati, and Marios Avraamides for making the space syntax data for Virtual Silcton freely available.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.cognition.2022.105360>.

References

- Allen, M., Poggiali, D., Whitaker, K., Rhys Marshall, T., van Langen, J., & Kievit, R. A. (2021). Raincloud plots: A multi-platform tool for robust data visualization. *Wellcome Open Research*, 4(63).
- Balaguer, J., Spiers, H., Hassabis, D., & Summerfield, C. (2016). Neural mechanisms of hierarchical planning in a virtual Subway network. *Neuron*, 90(4), 893–903. <https://doi.org/10.1016/j.neuron.2016.03.037>
- Barhorst-Cates, E. M., Meneghetti, C., Zhao, Y., Pazzaglia, F., & Creem-Regehr, S. H. (2021). Effects of home environment structure on navigation preference and performance: A comparison in Veneto, Italy and Utah, USA. *Journal of Environmental Psychology*, 74, Article 101580.
- Behrendt, S. (2014). *Lm.beta: Add standardized regression coefficients to lm-objects (1.5–1) [Computer Software]*.
- Bellmund, J. L. S., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*, 362(6415), eaat6766. <https://doi.org/10.1126/science.aat6766>
- Ben-Shachar, M., Lüdtke, D., & Makowski, D. (2020). Effectsize: Estimation of effect size indices and standardized parameters. *Journal of Open Source Software*, 5(56), 2815.
- Bohn, R., & Short, J. (2012). Measuring consumer information. *International Journal of Communication*, 6, 980–1000.
- Calhoun, A. J., & Hayden, B. Y. (2015). The foraging brain. *Current Opinion in Behavioral Sciences*, 5, 24–31. <https://doi.org/10.1016/j.cobeha.2015.07.003>
- Canty, A., & Ripley, B. (2021). *Boot: Bootstrap R (S-plus) functions (1.3-28) [Computer Software]*.
- Chentanez, N., Barto, A., & Singh, S. (2004). Intrinsically motivated reinforcement learning. *Advances in Neural Information Processing Systems*, 17.
- Chrastil, E. R., & Warren, W. H. (2012). Active and passive contributions to spatial learning. *Psychonomic Bulletin & Review*, 19(1), 1–23. <https://doi.org/10.3758/s13423-011-0182-x>
- Coutrot, A., Manley, E., Goodroe, S., Gahnstrom, C., Filomena, G., Yesiltepe, D., ... Spiers, H. J. (2022). Entropy of city street networks linked to future spatial navigation ability. *Nature*, 604, 104–110.
- Davidson, J. D., & El Hady, A. (2019). Foraging as an evidence accumulation process. *PLoS Computational Biology*, 15(7), Article e1007060.
- Davison, A. C., & Hinkley, D. V. (1997). *Bootstrap methods and their applications*. Cambridge University Press.
- De Cothi, W., Nyberg, N., Griesbauer, E. M., Ghanamé, C., Zisch, F., Lefort, J. M., ... Spiers, H. J. (2022). Predictive maps in rats and humans for spatial navigation. *Current Biology*, 32(17), 3676–3689.
- Devlin, S., Georgescu, R., Momennejad, I., Rzepecki, J., Zuniga, E., Costello, G., ... Hofmann, K. (2021, July). Navigation Turing Test (NTT): Learning to Evaluate Human-Like Navigation. In *International Conference on Machine Learning* (pp. 2644–2653). PMLR.
- Duncan, K. D., & Schlichting, M. L. (2018). Hippocampal representations as a function of time, subregion, and brain state. *Neurobiology of Learning and Memory*, 153, 40–56.
- Emo, B. (2014). Seeing the axial line: Evidence from wayfinding experiments. *Behavioral Science*, 4(3), 167–180.
- Epstein, R. A., Patai, E. Z., Julian, J. B., & Spiers, H. J. (2017). The cognitive map in humans: Spatial navigation and beyond. *Nature Neuroscience*, 20(11), 1504–1513. <https://doi.org/10.1038/nn.4656>
- Farran, E. K., Blades, M., Hudson, K. D., Sockeel, P., & Courbois, Y. (2022). Spatial exploration strategies in childhood; exploration behaviours are predictive of navigation success. *Cognitive Development*, 61, 101153.
- Freund, J., Brandmaier, A. M., Lewejohann, L., Kirste, I., Kritzler, M., Krüger, A., ... Kempermann, G. (2013). Emergence of individuality in genetically identical mice. *Science*, 340(6133), 756–759.
- Frick, A., Möhring, W., & Newcombe, N. S. (2014). Picturing perspectives: Development of perspective-taking abilities in 4-to 8-year-olds. *Frontiers in Psychology*, 5, 386.
- Furman, A. J., Clements-Stephens, A. M., Marchette, S. A., & Shelton, A. L. (2014). Persistent and stable biases in spatial learning mechanisms predict navigational style. *Cognitive, Affective, & Behavioral Neuroscience*, 14(4), 1375–1391.
- Gagnon, K. T., Cashdan, E. A., Stefanucci, J. K., & Creem-Regehr, S. H. (2016). Sex differences in exploration behavior and the relationship to harm avoidance. *Human Nature*, 27(1), 82–97. <https://doi.org/10.1007/s12110-015-9248-1>
- Gagnon, K. T., Thomas, B. J., Munion, A., Creem-Regehr, S. H., Cashdan, E. A., & Stefanucci, J. K. (2018). Not all those who wander are lost: Spatial exploration patterns and their relationship to gender and spatial memory. *Cognition*, 180, 108–117. <https://doi.org/10.1016/j.cognition.2018.06.020>
- Gardony, A. L., Taylor, H. A., & Brunyé, T. T. (2016). Gardony map drawing analyzer: Software for quantitative analysis of sketch maps. *Behavior Research Methods*, 48(1), 151–177.
- Ghosh, A., Puthusseryppady, V., Chan, D., Mascolo, C., & Hornberger, M. (2022). Machine learning detects altered spatial navigation features in outdoor behaviour of Alzheimer's disease patients. *Scientific Reports*, 12(1), 1–13.
- Gruber, M. J., & Ranganath, C. (2019). How curiosity enhances hippocampus-dependent memory: The prediction, appraisal, curiosity, and exploration (PACE) framework. *Trends in Cognitive Sciences*, 23(12), 1014–1025.
- Hegarty, M. (2002). Development of a self-report measure of environmental spatial ability. *Intelligence*, 30(5), 425–447. [https://doi.org/10.1016/S0160-2896\(02\)00116-2](https://doi.org/10.1016/S0160-2896(02)00116-2)
- Heller, A. S., Shi, T. C., Ezie, C. E. C., Reneau, T. R., Baez, L. M., Gibbons, C. J., & Hartley, C. A. (2020). Association between real-world experiential diversity and positive affect relates to hippocampal–striatal functional connectivity. *Nature Neuroscience*, 23(7), 800–804. <https://doi.org/10.1038/s41593-020-0636-4>
- Hillier, B., & Hanson, J. (1989). *The social logic of space*. Cambridge University Press.
- Howard, L. R., Javadi, A. H., Yu, Y., Mill, R. D., Morrison, L. C., Knight, R., ... Spiers, H. J. (2014). The Hippocampus and entorhinal cortex encode the path and Euclidean distances to goals during navigation. *Current Biology*, 24(12), 1331–1340. <https://doi.org/10.1016/j.cub.2014.05.001>
- Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3), 90–95.
- Ishikawa, T., & Montello, D. R. (2006). Spatial knowledge acquisition from direct experience in the environment: Individual differences in the development of metric knowledge and the integration of separately learned places. *Cognitive Psychology*, 52(2), 93–129. <https://doi.org/10.1016/j.cogpsych.2005.08.003>
- Javadi, A.-H., Emo, B., Howard, L. R., Zisch, F. E., Yu, Y., Knight, R., ... Spiers, H. J. (2017). Hippocampal and prefrontal processing of network topology to simulate the future. *Nature Communications*, 8, 14652. <https://doi.org/10.1038/ncomms14652>
- Julian, J. B., Keinath, A. T., Frazzetta, G., & Epstein, R. A. (2018). Human entorhinal cortex represents visual space using a boundary-anchored grid. *Nature Neuroscience*, 21(2), 191–194.
- Kaplan, R., Doeller, C. F., Barnes, G. R., Litvak, V., Düzel, E., Bandettini, P. A., & Burgess, N. (2012). Movement-related Theta rhythm in humans: Coordinating self-directed hippocampal learning. *PLoS Biology*, 10(2), Article e1001267. <https://doi.org/10.1371/journal.pbio.1001267>
- Karuza, E. A., Kahn, A. E., Thompson-Schill, S. L., & Bassett, D. S. (2017). Process reveals structure: How a network is traversed mediates expectations about its architecture. *Scientific Reports*, 7(1), 12733.
- Karuza, E. A., Thompson-Schill, S. L., & Bassett, D. S. (2016). Local patterns to global architectures: Influences of network topology on human learning. *Trends in Cognitive Sciences*, 20(8), 629–640. <https://doi.org/10.1016/j.tics.2016.06.003>
- Killian, N. J., Jutras, M. J., & Buffalo, E. A. (2012). A map of visual space in the primate entorhinal cortex. *Nature*, 491(7426), 761–764.
- Krupic, J., Bauza, M., Burton, S., Barry, C., & O'Keefe, J. (2015). Grid cell symmetry is shaped by environmental geometry. *Nature*, 518(7538), 232–235.
- Lydon-Staley, D. M., Zhou, D., Blevins, A. S., Zurn, P., & Bassett, D. S. (2021). Hunters, busybodies and the knowledge network building associated with deprivation curiosity. *Nature Human Behaviour*, 5(3), 327–336. <https://doi.org/10.1038/s41562-020-00985-7>
- Lynn, C. W., Kahn, A. E., Nyema, N., & Bassett, D. S. (2020). Abstract representations of events arise from mental errors in learning and memory. *Nature Communications*, 11(1), 2313. <https://doi.org/10.1038/s41467-020-15146-7>
- Momennejad, I. (2020). Learning structures: Predictive representations, replay, and generalization. *Current Opinion in Behavioral Sciences*, 32, 155–166. <https://doi.org/10.1016/j.cobeha.2020.02.017>
- Monaco, J. D., Rao, G., Roth, E. D., & Knierim, J. J. (2014). Attentive scanning behavior drives one-trial potentiation of hippocampal place fields. *Nature Neuroscience*, 17(5), 725–731.
- Munion, A. K., Stefanucci, J. K., Rovira, E., Squire, P., & Hendricks, M. (2019). Gender differences in spatial navigation: Characterizing wayfinding behaviors. *Psychonomic Bulletin & Review*, 26(6), 1933–1940. <https://doi.org/10.3758/s13423-019-01659-w>
- Nantais, M. (2019). Exploring cognitive maps through sketching. *Electronic Thesis and Dissertation Repository*. 6788. <https://ir.lib.uwo.ca/etd/6788>.

- Pagkratidou, M., Galati, A., & Avraamides, M. (2020). Do environmental characteristics predict spatial memory about unfamiliar environments? *Spatial Cognition and Computation*, 20(1), 1–32.
- Peer, M., Brunec, I. K., Newcombe, N. S., & Epstein, R. A. (2021). Structuring knowledge with cognitive maps and cognitive graphs. *Trends in Cognitive Sciences*, 25(1), 37–54. <https://doi.org/10.1016/j.tics.2020.10.004>
- Penn, A. (2003). Space syntax and spatial cognition: Or why the axial line? *Environment and Behavior*, 35(1), 30–65.
- Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497(7447), 74–79.
- Qian, W., Lynn, C. W., Klishin, A. A., Stiso, J., Christianson, N. H., & Bassett, D. S. (2022). Optimizing the human learnability of abstract network representations. *Proceedings of the National Academy of Sciences*, 119(35), Article e2121338119.
- R Core Team. (2018). *R: A language and environment for statistical computing (3.5.1) [Computer software]*. R Foundation for Statistical Computing. <https://www.R-project.org/>.
- Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature Neuroscience*, 16(4), 486–492. <https://doi.org/10.1038/nn.3331>
- Schapiro, A. C., Turk-Browne, N. B., Norman, K. A., & Botvinick, M. M. (2016). Statistical learning of temporal community structure in the hippocampus. *Hippocampus*, 26(1), 3–8. <https://doi.org/10.1002/hipo.22523>
- Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychological Science*, 30(11), 1561–1572.
- Schulz, L. (2012). The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in Cognitive Sciences*, 16(7), 382–389.
- Solway, A., Diuk, C., Córdova, N., Yee, D., Barto, A. G., Niv, Y., & Botvinick, M. M. (2014). Optimal behavioral hierarchy. *PLoS Computational Biology*, 10(8), Article e1003779. <https://doi.org/10.1371/journal.pcbi.1003779>
- Song, E. Y., Kim, Y. B., Kim, Y. H., & Jung, M. W. (2005). Role of active movement in place-specific firing of hippocampal neurons. *Hippocampus*, 15(1), 8–17. <https://doi.org/10.1002/hipo.20023>
- Spiers, H. J., & Maguire, E. A. (2006). Thoughts, behaviour, and brain dynamics during navigation in the real world. *NeuroImage*, 31(4), 1826–1840. <https://doi.org/10.1016/j.neuroimage.2006.01.037>
- Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, 20(11), 1643–1653. <https://doi.org/10.1038/nn.4650>
- Terrazas, A., Krause, M., Lipa, P., Gothard, K. M., Barnes, C. A., & McNaughton, B. L. (2005). Self-motion and the hippocampal spatial metric. *Journal of Neuroscience*, 25(35), 8085–8096. <https://doi.org/10.1523/JNEUROSCI.0693-05.2005>
- Tingley, D., Yamamoto, T., Hirose, K., Keele, L., & Imai, K. (2017). *Mediation: Causal mediation analysis (4.4.6) [computer software]*.
- Uttal, D. H., Meadow, N. G., Tipton, E., Hand, L. L., Alden, A. R., Warren, C., & Newcombe, N. S. (2013). The malleability of spatial skills: A meta-analysis of training studies. *Psychological Bulletin*, 139(2), 352–402. <https://doi.org/10.1037/a0028446>
- Vezzani, G., Gupta, A., Natale, L., & Abbeel, P. (2019). Learning latent state representation for speeding up exploration. *ArXiv Preprint (ArXiv:1905.12621)*.
- Walter, J. L., Essmann, L., König, S. U., & König, P. (2022). Finding landmarks—an investigation of viewing behavior during spatial navigation in VR using a graph-theoretical analysis approach. *PLoS Computational Biology*, 18(6), Article e1009485.
- Weisberg, S. M., & Newcombe, N. S. (2016). How do (some) people make a cognitive map? Routes, places, and working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(5), 768.
- Weisberg, S. M., Schinazi, V. R., Ferrario, A., & Newcombe, N. S. (2022). Evaluating the effects of a programming error on a virtual environment measure of spatial navigation behavior. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Weisberg, S. M., Schinazi, V. R., Newcombe, N. S., Shipley, T. F., & Epstein, R. A. (2014). Variations in cognitive maps: Understanding individual differences in navigation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(3), 669.
- Wiener, J. M., & Mallot, H. A. (2003). 'Fine-to-coarse' route planning and navigation in regionalized environments. *Spatial Cognition and Computation*, 3(4), 331–358.
- Wiener, J. M., Schnee, A., & Mallot, H. A. (2004). Use and interaction of navigation strategies in regionalized environments. *Journal of Environmental Psychology*, 24(4), 475–493.
- Winter, S. S., Mehlman, M. L., Clark, B. J., & Taube, J. S. (2015). Passive transport disrupts grid signals in the Parahippocampal cortex. *Current Biology*, 25(19), 2493–2502. <https://doi.org/10.1016/j.cub.2015.08.034>
- Wolbers, T., & Hegarty, M. (2010). What determines our navigational abilities? *Trends in Cognitive Sciences*, 14(3), 138–146. <https://doi.org/10.1016/j.tics.2010.01.001>
- Wolbers, T., & Wiener, J. M. (2014). Challenges for identifying the neural mechanisms that support spatial navigation: The impact of spatial scale. *Frontiers in Human Neuroscience*, 8, 571. <https://doi.org/10.3389/fnhum.2014.00571>
- Wu, C. M., Schulz, E., Garvert, M. M., Meder, B., & Schuck, N. W. (2020). Similarities and differences in spatial and non-spatial cognitive maps. *PLoS Computational Biology*, 16(9), Article e1008149. <https://doi.org/10.1371/journal.pcbi.1008149>
- Zurn, P., & Bassett, D. S. (2018). On curiosity: A fundamental aspect of personality, a practice of network growth. *Personality Neuroscience*, 1. <https://doi.org/10.1017/pen.2018.3>