The format of the cognitive map depends on the structure of the environment

Michael Peer*, Catherine Nadar*, Russell A Epstein

Department of Psychology, University of Pennsylvania, 3710 Hamilton Walk, Philadelphia, PA 19104, USA

* Equal contribution

Correspondence: mpeer@sas.upenn.edu

Humans and animals form cognitive maps that allow them to navigate through large-scale environments. Despite decades of research on these maps, a central question remains unclear: are these maps similar in nature across all environments, or are different kinds of maps formed in different kinds of environments? To investigate this, we examined spatial learning within three virtual environments: an open courtyard with patios connected by paths (open maze), a set of rooms connected by corridors (closed maze), and a set of isolated rooms connected only by teleporters (teleport maze). Importantly, all three environments shared the same topological graph structure. Post-learning tests showed that the environmental structure affected the accuracy, format, and variability of participants' spatial representations. The open maze was the most accurately remembered, followed by the closed maze, and then the teleport maze. Both Euclidean and graph-like spatial codes were formed in the open and closed mazes, but participants' navigational trajectories were more biased by graph knowledge (connectivity between rooms) in the closed maze compared to the open maze. Finally, performance in the open maze and teleport maze were relatively homogenous across participants, whereas performance in the closed maze exhibited greater individual variability. These results indicate that the structure of the environment strongly shapes the nature of the spatial representations that are formed within that environment, and that experimental findings obtained in any single environment may not generalize to others with different structure.

Introduction

How does the mind represent large-scale, navigable spaces? Decades of research focusing on this question have led to intense debate about the underlying nature of spatial knowledge, with some researchers arguing that it takes the form of a Euclidean map (Gallistel, 1990; O'Keefe & Nadel, 1978) and others arguing that it takes a more graph-like form (Kuipers, 1982; Warren, 2019). However, despite this extensive previous work, a crucial factor is often ignored – variability of structure across the spaces being represented. Most studies employ a single type of environment, such as an open arena, or a closed-in maze, and different environmental types are rarely compared within the same study. Although the implicit assumption is often made that findings obtained in one environment should generalize to all, this assumption has not been tested, and there are several reasons to believe that people may in fact form different kinds of mental representations in different kinds of environments.

One line of evidence comes from electrophysiological studies in rodents. When animals navigate in open arenas, hippocampal place cells typically fire in consistent locations irrespective of the direction from which these locations are accessed, indicating a direction-invariant place code (Moser et al., 2008; O'Keefe & Dostrovsky, 1971). In contrast, when animals navigate in linear corridors, place cell firing is modulated by the direction of movement, indicating a direction-dependent representation (McNaughton et al., 1983; Mehta et al., 1997; Muller et al., 1994). Another line of evidence comes from human cognitive studies. People navigating in large open environments exhibit behaviors and neural patterns that are consistent with the use of Euclidean cognitive maps (Chadwick et al., 2015; Doeller et al., 2010; Jacobs et al., 2013; Maidenbaum et al., 2018; Shine et al., 2019), whereas people navigating in maze-like environments exhibit behaviors and neural patterns that are consistent with the use of spatial representations that do not obey Euclidean rules (Chrastil & Warren, 2014; Doner et al., 2022; Ericson & Warren, 2020; He & Brown, 2019; Moeser, 1988; Muryy & Glennerster, 2018; Zetzsche et al., 2009). Taken as a whole, these literatures suggest that the structure of the environment may matter: whereas some environments might be more easily encoded using a Euclidean reference frame, others might be more easily encoded using a cognitive graph consisting of place nodes and their connections (Peer et al., 2021).

If people do tend to form different representations in different environments, this tendency might interact in interesting ways with individual differences in navigational ability. Previous work has shown that people differ in their capacity to form spatial representations of large-scale spaces (Ishikawa & Montello, 2006; Weisberg & Newcombe, 2018). For example, one study found that people learning a virtual environment could be separated into three groups: integrators, who could accurately point to locations in separately-learned parts of the environment; non-integrators, who could only point to locations in the same part of the environment but not across parts; and imprecise navigators, who could not point accurately to any location (Weisberg et al., 2014). It is unclear how these groupings, obtained in an environment consisting of bounded routes, would generalize to other types of environments such as large open spaces. One possibility is that individual differences might manifest differently across environments with different structure. Alternatively, these individual groupings might be stable character traits, perhaps shaped by the environment that the participants grew up in (Barhorst-Cates et al., 2021; Coutrot et al., 2022).

Here we set out to test these ideas, by investigating how large-scale spatial knowledge (often referred to as "cognitive maps") differs across newly learned environments with different spatial structure. Participants were randomly assigned to learn one of three virtual environments (Figure 1). The environments all consisted of nine connected subspaces with the same topological graph structure, but differed in terms of the cues available for determining location and heading in Euclidean space. In the first environment ("open maze"), the subspaces were unwalled patios connected by walkways, contained within a courtyard with distal landmarks. In the second environment ("closed maze"), the subspaces were enclosed rooms connected by corridors. In the third environment ("teleport maze"), the subspaces were enclosed rooms that were connected by "teleporters" located at the room centers. Following learning, participants performed a series of behavioral tests designed to explore the structure of their spatial representations (Figure S1). We hypothesized that the difference between the environments would affect both the quality and the format of the resultant spatial codes. To anticipate, we found that the environmental structure affected participants' ability to learn the environment, the format of their mental representation (Euclidean space or cognitive graph), and how these representations varied between participants.

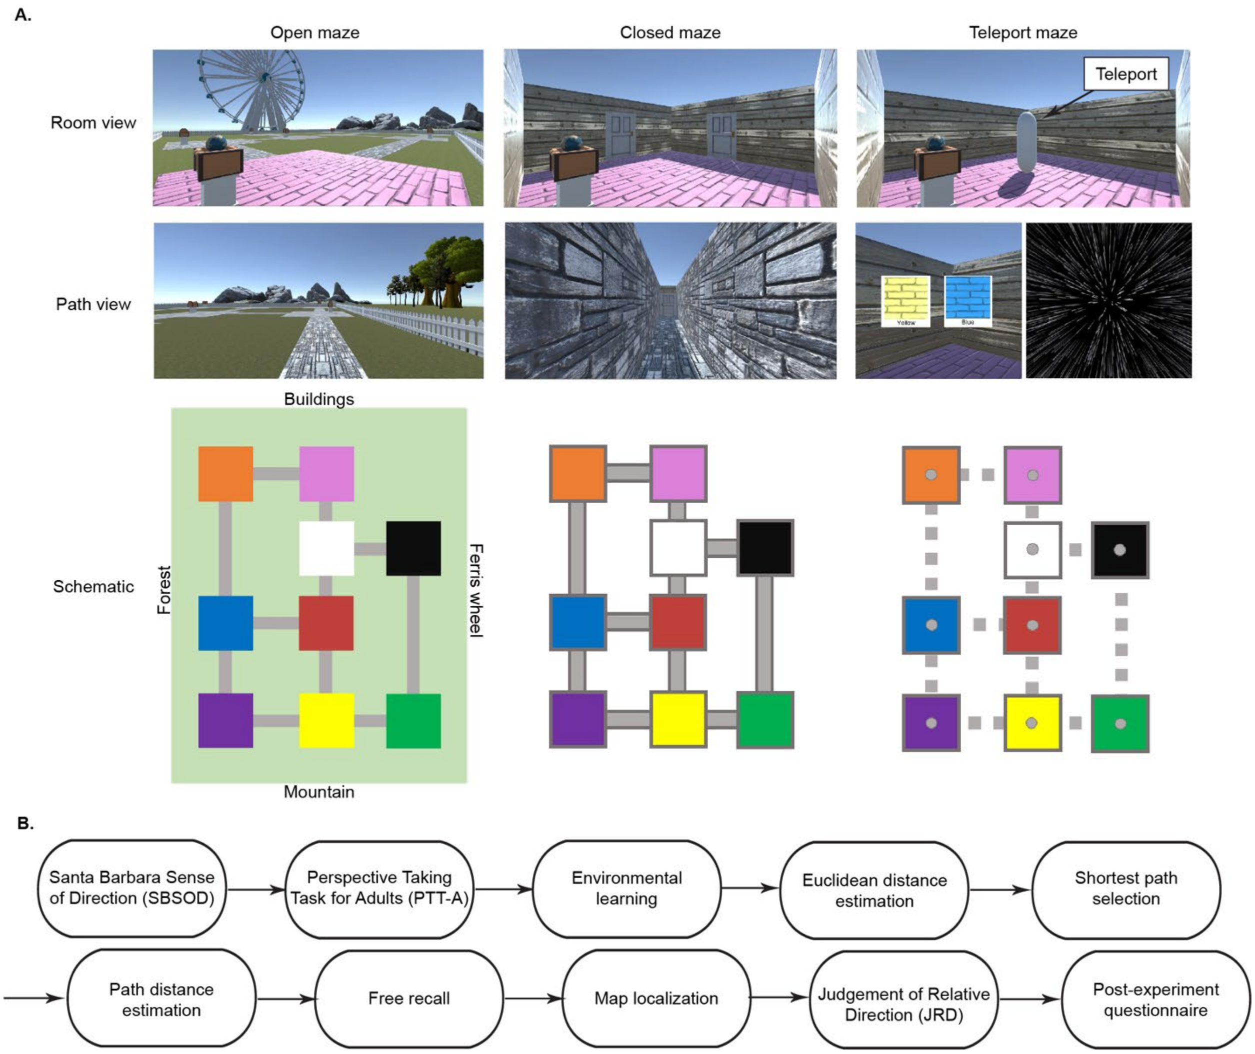

Figure 1: Experimental environments and procedure. A) Experimental environments. Participants learned to navigate in a virtual environment containing nine rooms connected by pathways, each with an associated floor color and object. In the open-maze condition (left), there were no walls and the entire environment and its surrounding distal landmarks were visible. In the closed-maze condition (center), the rooms and pathways were surrounded by walls so that only the immediate surroundings were visible, and participants had to maintain an internal sense of direction and location to navigate. In the teleport maze condition (right), rooms were physically isolated from each other, and participants relied on a system of teleporters to transition between rooms; therefore, they could not maintain a consistent sense of direction and could only learn the connectivity structure of the environment. Each environment is depicted from two viewpoints: Room view, which shows a part of the environment from first person point of view, and Path view, which shows a part of the environment from the corridor or teleport. B) Experimental procedure. Following learning, participants' memory of the environment was tested using a battery of spatial memory tasks. All participants performed all tasks consecutively, except for the teleport maze condition participants who did not do the Judgment of Relative Direction (JRD) task.

Results

The structure of the environment affects the accuracy of spatial knowledge

Each participant was familiarized with one of the three virtual environments through a multi-stage training procedure. Participants first learned to navigate between the rooms (defined by their floor colors), and then learned to navigate to objects located in “treasure chests” within each room. By the end of the training, participants were able to navigate to all rooms and objects without error, even when the identifying floor colors of the rooms were not shown, and the objects remained hidden in the treasure chests. In addition, participants’ routes to the targets were significantly shorter than a random walk ($p<0.01$ for all three environments), indicating that they formed a representation of the environmental structure instead of walking randomly until they found their target.

Participants then performed a series of memory tasks to assess their spatial knowledge of the environment they learned. Three tasks were designed to test Euclidean knowledge (Euclidean distance estimation, map localization, and judgment of relative direction) and two were designed to test graph-based knowledge (path distance estimation, shortest path selection). All tasks were administered to all participants, with the exception the judgment of relative direction task (JRD), which was not administered to participants in the teleporter maze because directions between rooms is not a knowable quantity in this environment. In all three environments, performance in all tasks was better than chance (all $ps<0.01$; Figure 2A-E). However, we also observed consistent differences in performance between the environments: accuracy was highest in the open maze condition, intermediate in the closed maze condition, and lowest in the teleport maze condition. These differences between environments were confirmed for the Euclidean distance estimation, map localization, path distance estimation, and shortest path selection tasks by ANOVA (all $Fs>18$, all $ps<0.0001$; all p values FDR-corrected across tasks) and Tukey-Cramer post-hoc tests of pairwise differences between environments (all $ps<0.05$, except for the difference between the closed maze and teleport maze in the Euclidean and path distance estimation tasks: $p=0.0504$, 0.074, respectively). The difference between environments was confirmed in the JRD task by t-test between open and closed maze performance ($t(37)=2.34$, $p=0.025$). Notably, the same pattern of performance across environments was observed for the three tasks that assessed Euclidean knowledge (Euclidean distance estimation, map localization, JRD) and the two tasks that assessed graph-based knowledge (path distance estimation, shortest path selection).

We also administered a free recall task on the object names. Previous work has shown that the order of free recall can be shaped by the spatial proximity of objects, such that closer items are more likely to be recalled in sequence (Hirtle & Jonides, 1985; Miller et al., 2013; Peer & Epstein, 2021). Consistent with this prior work, the order of object recall was related to the both the Euclidean and path distance between objects in the open maze (Euclidean distance between recalled objects vs. random recall distances: $t(19)=3.84$, $p=0.003$; Path distance: $t(19)=3.68$, $p=0.005$; all ps FDR-corrected across environments) and closed maze (Euclidean distance: $t(19)=2.68$, $p=0.02$; path distance: $t(19)=2.74$, $p=0.02$). In the teleport maze, recall order was related to path distance ($t(19)=2.99$, $p=0.01$) but there was no relation between recall order and the Euclidean distance ($t(19)=-0.16$, $p=0.87$). The last finding is not surprising given that participants did not learn anything about the “true” spatial positions of the rooms in the teleport condition.

Taken together, these results demonstrate that participants can acquire information on the structure of different environments, even if the environments only contain information about connectivity structure (in the teleport maze condition). However, there is an ordering of spatial memory accuracy that relates to the richness of the location and direction cues. Accuracy is highest in the open maze, where participants can determine their location and heading directly through perception. Accuracy is intermediate in the closed maze, where participants can use perception to determine their local (within-room) location and heading, but can only determine their global (across-rooms) location and heading by keeping track of these quantities as they move from room to room. Accuracy is lowest in the teleport maze, where any inference about global location and heading must be made indirectly, because there were no cues about heading or location as the participants teleported from room to room. Thus, the structure of the environment affects the quality of the spatial representations formed.

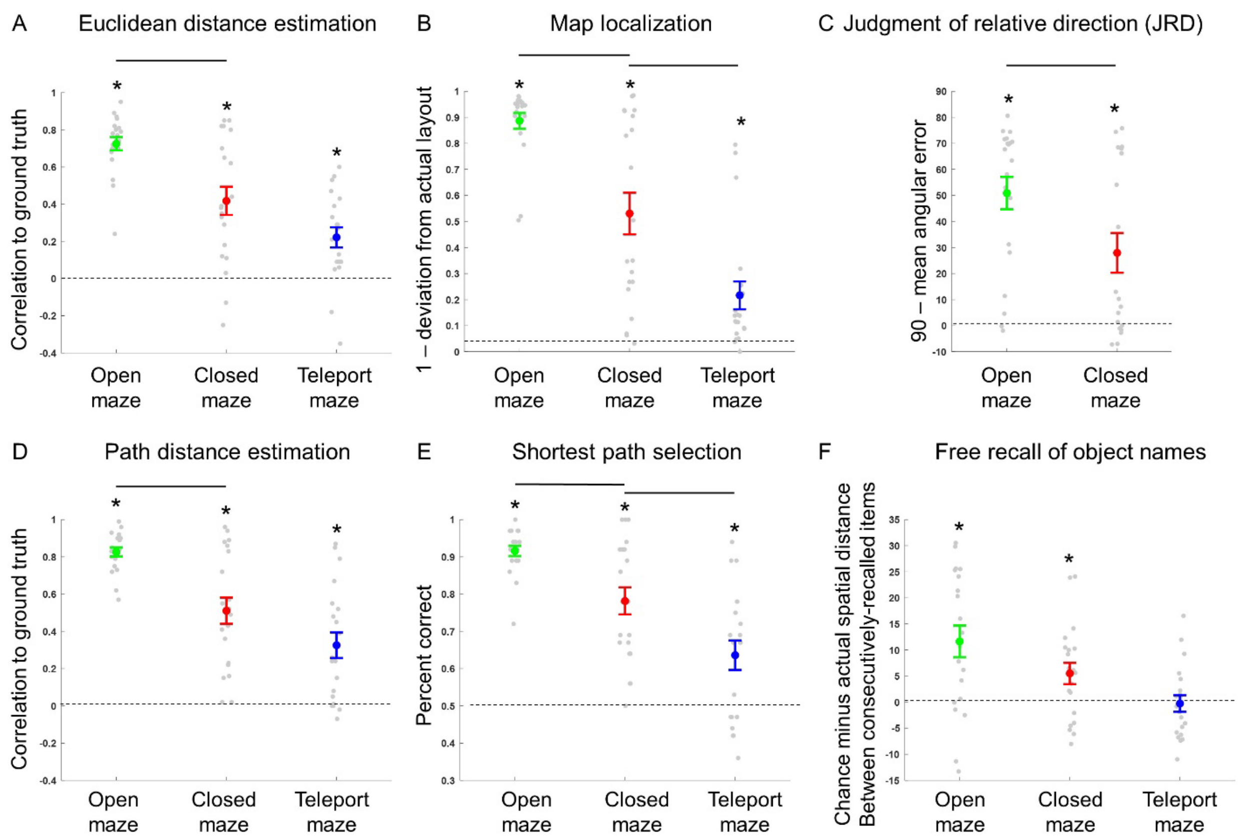

Figure 2: Participants' spatial memory varies across environments. A-E) Performance in spatial memory tasks. In all environments and tasks, participants performed above chance level, indicating that they acquired information on the environment's spatial layout. However, spatial memory accuracy across tasks was highest in the open maze condition, lower in the closed maze condition, and lowest in the teleport maze condition, indicating that the gradual removal of location and direction cues impaired spatial memory despite the similar layout and connectivity of all environments. F) Free recall of object names: objects' recall order was related to inter-object spatial distances in the open and closed maze conditions, but not in the teleport maze condition. Grey dots represent individual participants' data; Colored dots indicate group means (green – open maze, red – closed maze, blue – teleport maze); error bars indicate standard error; asterisks represent significant above-chance performance; full lines represent significant differences between adjoining conditions; dashed lines represent chance level.

The structure of the environment affects the format of spatial knowledge

We next investigated whether the structure of the environment affects not just the accuracy of learned spatial representations, but also their format. We were specifically interested in the distinction between map-like and graph-like representations. Some researchers have argued that spatial knowledge typically takes the form of a Euclidean map-like code (Gallistel, 1990; O'Keefe & Nadel, 1978; Siegel & White, 1975), while others have argued that it typically takes the form of a graph consisting of nodes and edges (Kuipers, 1982; Warren, 2019). We previously suggested that Euclidean representations might be more common in open environments, while graph-like representation might be more common in closed environments (Peer et al., 2021). Here we tested this idea.

To do this, we first examined route preferences during learning, focusing on trials in the open and closed mazes that had two possible routes toward the target, each with the same Euclidean path length but a different number of intervening rooms (Figure 3A). We reasoned that if participants relied on a purely Euclidean cognitive map, then Euclidean distance to goal should be the only consideration for route choice, and both routes should be selected with equal probability. If, however, participants relied on a graph-like representation, then they should prefer to take the path with fewer intervening rooms, because this involves travelling over fewer graph edges. Moreover, if graph-like representations are more common in closed environments, this bias would be stronger in the closed maze than in the open maze. Consistent with this prediction, participants in the closed maze preferred routes passing through a smaller number of rooms ($p=0.001$), while participants in the open-maze participants did not ($p=0.17$, FDR-corrected across environments), and the difference between the mazes was significant ($t(38)=2.11$, $p=0.04$, two-samples two-tailed t-test; Figure 3B).

Next we tested for analogous effects in the memory tasks. First, we looked at route preferences in the shortest-path selection task. When asked to choose between two paths that had equal Euclidean length but different numbers of path segments, both open- and closed-maze participants consistently selected routes with fewer intervening rooms (both $ps<0.0001$, FDR-corrected for multiple comparisons), indicating the use of a graph-like representation. However, in this case, the tendency to use graph knowledge did not significantly differ between the two environments ($t(38)=0.95$, $p=0.35$, two-samples two-tailed t-test; Figure 3C). Second, we looked at distance estimates in the Euclidean distance estimation task. We reasoned that if participants utilized a graph-like code instead of (or in addition to) a map-like code, then their responses when estimating Euclidean distance should be affected by the path distance between the two items on each trial, even though path distance is not what the task queries. In the open maze, correlation of Euclidean distance estimates to the veridical Euclidean distances (mean $r=0.72$) was equivalent to correlation to the veridical path distances (mean $r=0.74$, $p=0.40$ for the difference; Figure 3D). By contrast, in the closed maze, correlation of Euclidean distance estimates to veridical Euclidean distances (mean $r=0.42$) was significantly lower than correlation to path distance ($r=0.46$, $p=0.005$ for the difference; Figure 3D). However, the difference between the Euclidean and path correlation was not significantly larger in the closed maze compared to the open maze ($p=0.59$).

In sum, our results suggest that participants' spatial knowledge is not restricted to a Euclidean map, but includes graph-like elements. We also found some evidence that participants were more likely to use a graph-like code in closed environments than in open environments. However, the evidence for this tendency was mixed: it was observed when participants were actually navigating in the environment (in the learning task), but it was not observed in the post-navigation memory tasks.

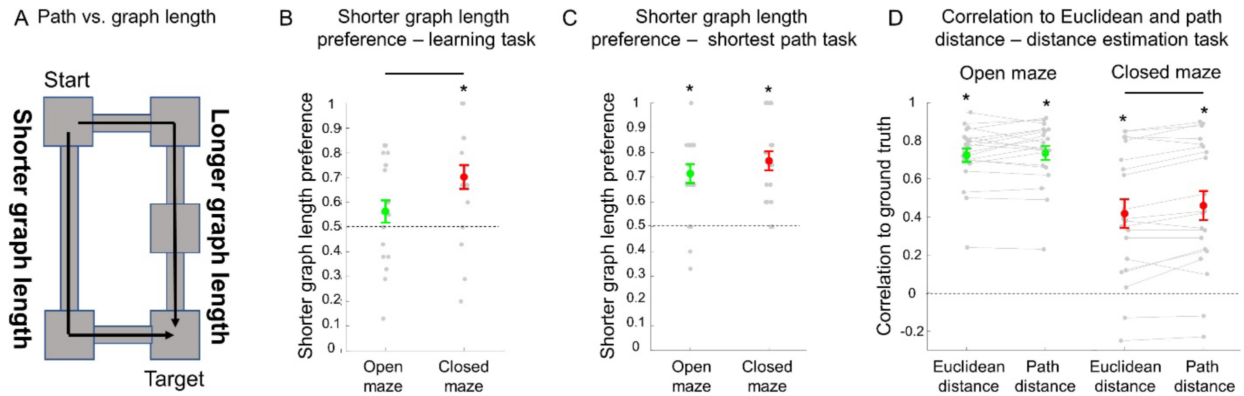

Figure 3: Evidence for use of cognitive graph knowledge. A) A schematic of a sample trial, where there were two pathways between the starting and target objects with equal path length but different number of intervening rooms (graph length). If participants form a Euclidean cognitive map and use it to navigate, both pathways should be selected with equal probability; but if participants form a cognitive-graph based representation, they might choose the route with the lesser number of intervening rooms (shorter graph length). B) Route selection during the learning task: participants were biased toward selection of the route with the smaller number of intervening rooms in the closed maze condition but not in the open maze condition. C) Route selection during the shortest path selection task: participants were biased toward selection of the route with the smaller number of intervening rooms in both environments. D) In the Euclidean distance estimation task, participants in the closed maze condition had a higher correlation between their estimates and the path distances between locations than to the Euclidean distances between locations, indicating a distortion of perceived Euclidean distances by cognitive graph knowledge. Plot elements are the same as in Figure 2.

The structure of the environment affects how spatial knowledge varies across individuals

Finally we turned to an examination of how spatial knowledge varied across individual participants. Previous work has examined individual differences in cognitive maps within a single environment (Ishikawa & Montello, 2006; Weisberg & Newcombe, 2018), and we wanted to understand how this variability might express itself across different types of environments. To investigate this issue, we focused first on the map localization data, as this is the task that most directly assesses participants' ability to form an accurate cognitive map of the environment.

Inspection of the data (Fig 4A) revealed an interesting pattern of variability across the three mazes. In the open maze, most participants were able to localize the objects and rooms with high accuracy, but there were two participants who exhibited notably worse performance. In the teleport maze, the converse pattern was observed: most participants were unable to localize the items accurately, but there were three participants with notably better performance. This suggests that some teleport maze participants were able to intuit an accurate cognitive map of the environment even in the absence any direct cues about the relative locations of the rooms. Finally, participants in the closed maze showed a wide range of performance, with some participants responding as accurately as the majority of the open maze participants, and some responding as inaccurately as the majority of the teleport maze participants (Figure S2).

The overall picture suggests a division of participants into two groups: participants who can form an accurate cognitive map, as evidenced by their ability to localize items accurately in the map test, and those who cannot. This division was confirmed by a clustering analysis (clustering silhouette score = 0.91, $p<0.001$). Following the terminology used in Weisberg & Newcombe (Weisberg & Newcombe, 2018), we refer to these two groups as “integrators” (accurate localization) and “non-integrators” (inaccurate localization). These groups were unequally distributed across the three mazes: in the open maze, most participants were integrators; in the closed maze, about half of the participants were integrators and half were non-integrators; in the teleport maze, most participants were non-integrators. The proportion of integrators to non-integrators differed significantly between the mazes ($p<0.01$, Chi-square test). These data suggest that a random participant’s propensity to behave as an integrator or non-integrator depends on the structure of the environment.

We next examined the variability of individual performance in the other memory tasks. We were particularly interested in whether the distinction between integrators and non-integrators, derived from map localization performance, would be maintained in these tasks. For the statistical analyses, we divided the participants into four groups: open maze integrators, closed maze integrators, closed maze non-integrators, and teleport maze non-integrators. We did not analyze open maze non-integrators and teleport maze integrators because there were only small numbers of participants in these groups.

Performance differed across participant groups in the Euclidean Distance Estimation Task ($F(3,54)=37.2$, $p<0.0001$), the path distance estimation task ($F(3,54)= 32.6$, $p<0.0001$) and the shortest path selection task ($F(3,54)=41.8$, $p<0.0001$; Fig. 4B-E). Post-hoc pairwise comparisons revealed an interesting pattern that was common for all three tasks: performance of closed-maze integrators was significantly better than the performance of closed-maze non-integrators ($p<0.0001$ in all tasks), but performance did not significantly differ between open-maze integrators and closed-maze integrators ($p>0.91$ in all tasks), nor did it significantly differ between closed maze non-integrators and teleport maze non-integrators ($p>0.55$ in all tasks). In other words, integrators in the open and closed mazes performed similarly, and non-integrators in the closed and teleport mazes performed similarly. A similar division between integrators and non-integrators was observed in the JRD task. Performance in this task differed across open maze integrators, closed maze integrators, and closed maze non-integrators ($F(2,36)=23.12$, $p<0.0001$; Fig. 4C). Post-hoc pairwise comparison found that performance differed between closed-maze integrators and closed maze non-integrators ($p<0.0001$), but not between open-maze integrators and closed-maze integrators ($p=0.86$).

What drives the division of participants into integrators and non-integrators? Our findings suggest that the environment itself is a primary factor: most participants in the open maze are integrators, whereas most participants in the teleport maze are non-integrators. However, the control of the environment is not absolute, and individual differences are observed in all three environments. We tested whether this individual variability relates to participants’ baseline spatial abilities, by examining the correlation between participants’ performance in each task and their perspective-taking ability (PTTA score) and self-rated navigational ability (SBSOD score). To control for the difference between conditions, we analyzed data from each maze separately (Table S1). Participants’ perspective taking ability was found to be correlated to their map localization performance in the open maze environment condition ($r=0.67,p<0.0001$) and to their path distance estimation, JRD performance and learning task accuracy in the closed maze condition ($r=0.55,0.57,0.54$, $p=0.043,0.043,0.043$, respectively; FDR-corrected across tasks and environmental conditions for SBSOD and PTTA separately). We did not observe any relationship between task performance and SBSOD scores.

Overall, our results suggest that individual variability in environmental learning is related to participants’ baseline perspective taking ability. In addition, this variability is not expressed the same way across different environments; participants almost uniformly perform well in open environments, almost uniformly perform poorly in cue-limited (teleport maze) environments, and the individual variability between them becomes most pronounced in maze-like environments (which have limited visibility and require the internal tracking of location and direction).

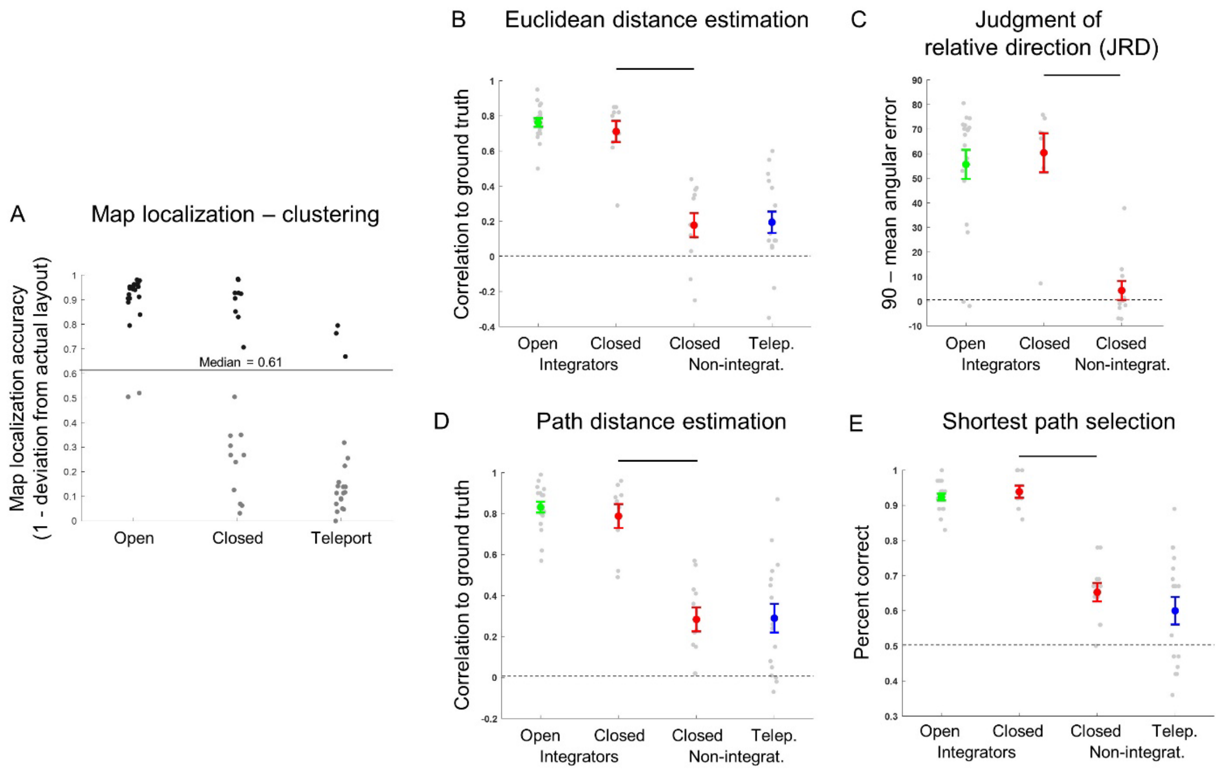

Figure 4: Individual variability is manifested differently in different environments. (A) Dot plot of map localization task performance, with clustering to two groups ("integrators" and "non-integrators") according to median performance. The proportion of integrators varies across environments (almost all participants in the open maze, about half in the closed maze, and almost no participants in the teleport maze), suggesting that the ability to integrate is not a stable personality trait but instead depends on the environment being navigated. (B-E) Spatial memory task performance after dividing participants into "integrators" and "non-integrators" (excluding the small groups of open maze non-integrators and teleport maze integrators). Across tasks, performance among the closed maze integrators was similar to open maze integrators, and performance among the closed maze non-integrators was similar to the teleport maze non-integrators. The open maze and closed maze integrators consistently performed better than the teleport maze and closed maze non-integrators. These results indicate that the environment plays a role in participants' ability to form an accurate cognitive map. Plot elements are the same as in Figure 2.

Discussion

The aim of this study was to understand how the physical structure of the navigable environment affects the accuracy and format of cognitive maps. Our analyses yielded three main findings. First, the accuracy of spatial knowledge was affected by the structure of the environment. Despite the fact that the open maze, closed maze, and teleport maze all had the same number of rooms and the same connectivity structure, spatial memory differed substantially between these conditions, with greater accuracy in environments with greater number of orienting cues (open>closed>teleport). Second, the format of spatial knowledge was also affected by the structure of the environment. Although we found evidence for both map-like and graph-like representations, participants used graph-like codes to guide navigation more often in the closed maze than in the open maze. Third, inter-individual variability in spatial knowledge was affected by the structure of the environment. Variability was relatively low in the open maze, where most participants were classified as integrators, and it was also low in the teleport maze, where most participant were classified as non-integrators. In contrast, variability was higher in the closed maze, where participants fell about evenly into these two groups. Notably, however, there were good and bad performers in all environments, suggesting that some individual differences may be impervious to environmental structure. Taken together, these results emphasize the strong effect that environmental structure has on spatial knowledge, and indicate that past and future studies of spatial representations should be carefully interpreted with respect to the specific type of environment used in each study. Below we discuss each of our main findings in more detail.

The first question we asked was whether the structure of the environment affects the accuracy of spatial knowledge. We found strong evidence that it does. There was a consistent ordering of performance across the environments: accuracy in all of our spatial memory tasks was highest for participants in the open maze, intermediate for participants in the closed maze, and lowest for participants in the teleport maze. This was despite the fact that all three environments had the exact same number of rooms, room geometry, and connectivity structure.

What could be the cause of this ordering? By design, the three environments differed in the cues available for determining heading and location in global Euclidean space. In the open maze, the entire environment was always visible, including the boundary of the courtyard and the distal landmarks. Thus, it was possible for participants to determine their global location and heading directly through perception. In the closed maze, on the other hand, participants could not see beyond the walls that bounded the local room or corridor. Thus, they could determine their local (within-room) location and heading through direct perception, but their global position and heading could only be ascertained by using an active memory process to keep track of these quantities as they moved from room to room. In the teleport maze, even this active tracking process was not possible, because there were no distance cues when participants were teleported from room to room, and their heading when placed in the new room was randomly chosen. Thus, the key factor that likely accounts for the difference in performance across the three environments is the ability of the participants to maintain a sense of global location and direction in Euclidean space during learning. This conclusion is consistent with a large literature that suggests that the ability to path integrate—to keep track of one's position and heading during exploration—is crucial for building up a cognitive map (Etienne & Jeffery, 2004; McNaughton et al., 1996, 2006).

The second question that we asked was whether the format of the environment affects the format of the cognitive map. The question of format has been the subject of considerable ongoing debate. Some theories posit that spatial knowledge consists of a map of locations represented within a global Euclidean coordinate system (Gallistel, 1990; O'Keefe & Nadel, 1978; Siegel & White, 1975), while others theories posit that it consists of a graph of routes connecting different locations, with no integration of the locations into a single global reference frame (Kuipers, 1982; Meilinger, 2008; Warren, 2019). In a previous review, we argued that map-like and graph-like representations might simultaneously exist in the same individuals and be differentially employed in different environments (Peer et al., 2021). Our findings provide evidence in support of these ideas.

Consideration of results across the three environments suggests that people formed both Euclidean and graph-like representations. Supporting the use of Euclidean codes is the fact that participants were able to perform tasks that were designed to tap Euclidean knowledge (map localization, Euclidean distance estimation, JRD). Moreover—as discussed in the previous section—their ability to do so varied across environments depending on the presence or absence of cues that allowed them to determine their global Euclidean position and heading. Of particular note, the difference in performance between the open and closed maze would not be found if only graph-like representations were formed, because such representations only encode local spatial features (node identity, angle of links at a node) that are equally present in both the open and closed maze (Ericson & Warren, 2020; Warren, 2019). However, we also found evidence for graph-like representations. Participants' navigational choices during learning and their responses in the shortest path task were influenced by graph knowledge: when faced with a decision between two paths of equal Euclidean length in the open and closed mazes, they preferred the path with fewer graph segments. Moreover, in the teleport maze, people were able to use pure graph knowledge (connectivity between locations) to navigate, albeit with less accuracy, suggesting that spatial representations can, in some circumstances, be predominantly graph-based. Thus, neither Euclidean maps nor cognitive graphs can explain our results on their own; both are needed.

We initially hypothesized that the use of graph knowledge would differ across environments, and we found some evidence in support of this hypothesis. When navigating through the environment in the learning phase, participants' preference to choose paths with fewer graph segments (when faced with alternatives of equal length) was greater in the closed maze compared to the open maze. This preference is consistent with our previous suggestion that the difficulty of integrating subspaces into a global Euclidean representation in a closed environment might lead to greater reliance on environmental topology (Peer et al., 2021). However, this effect was not observed when participants made route choices in the shortest path selection task. Thus, the evidence that participants form representations that are more graph-like when they are in closed environments is mixed.

The final question that we asked was how individual variability would manifest across the three environments. Previous reports have demonstrated marked differences between people in their ability to integrate subparts of the environment to form a global cognitive map (e.g. (Ishikawa & Montello, 2006; Weisberg et al., 2014; Weisberg & Newcombe, 2018)). For example, in an environment consisting of constrained routes, people can be separated into: integrators, who can accurately point to locations in separately-learned parts of the environment; non-integrators, who can point to locations in the same part of the environment but cannot point across parts; and imprecise navigators, who cannot accurately point to any location (Ishikawa & Montello, 2006; Weisberg et al., 2014; Weisberg & Newcombe, 2018). These differences have been shown to relate both to general cognitive abilities like mental rotation and perspective taking (Weisberg et al., 2014), as well as to the type of environment people grew up in (e.g. the layout of streets in people's home town (Barhorst-Cates et al., 2021; Coutrot et al., 2022)). However, these studies all used a single environment, leaving open the question of whether the observed individual differences are stable character traits, or tendencies that might manifest differently in the same individual depending on environmental context.

Here we found that there was a strong interaction between individuals' cognitive mapping ability and the structure of the environment: in the open maze most participants were integrators, in the teleport maze most participants were non-integrators, and in the closed maze participants fell about equally into the two groups. But there were some notable exceptions: some participants performed poorly in the open maze despite the abundance of spatial cues, and some participants performed well in the teleport maze (even in the Euclidean tasks), despite a paucity of cues. These results suggest that cognitive mapping ability manifests differently across environments in most people, but there are some individuals on the top or bottom end of the spectrum who show a greater degree of stability.

Previous studies of individual differences in cognitive mapping have used environments that allowed continuous tracking of location and heading, but without full visibility to other parts of the environment (Ishikawa & Montello, 2006; Weisberg & Newcombe, 2018). The variability observed in these studies might be less robust in more open environments, or in graph-like environments like our teleport maze. Furthermore, the ability to mentally track location and heading might depend on perspective taking ability, which we found to be correlated to the individual differences in our task, in line with previous studies (Allen et al., 1996; Fields & Shelton, 2006; Kozhevnikov et al., 2006; Schinazi et al., 2013). Overall, this set of findings has two major implications: first, past (and future) studies of variability in navigational ability should be carefully interpreted with respect to the specific environment being explored; and second, specific environmental features can reduce the variability between people and allow some “bad navigators” to navigate efficiently and integrate between different environmental subparts, suggesting implications for real-life environmental design.

Our study has several limitations. First, we used environments that differed from each other along more than one feature (e.g. the open maze differed from the closed maze both in visibility across the environment and in the existence of distal landmarks). Further studies might attempt to disentangle the effect of each of these features on the resulting spatial representations, and test how the observed effects generalize to other environments. Second, our study did not investigate within-subject effects – although our individual difference analyses suggest that some individuals would perform differently in different environments, we do not have direct evidence that this is the case. Third, the study population was predominantly young adults (many of whom were university students), and participants were educated members of western societies who are familiar with built urban environments (e.g. cities with rectilinear streets, buildings with corridors and rooms). Further studies are needed to confirm the generality of the findings for other age groups and population sectors, and for human populations whose primary experience is with natural environments. Finally, some features of the environments (e.g. teleportation in the teleport maze, and virtual desktop navigation) are not realistic; future studies may study how these effects are manifested in real-life environments.

In conclusion, we found that specific features of each environment affect the accuracy and format of the mental representations people form of these environments. We further found that individual variability in cognitive map formation and use of graph knowledge are not constant but instead depend on the environment being learned. These findings suggest that care should be taken to consider the specific environment’s features when interpreting the spatial navigation literature, and that environments and navigational aids can be designed to facilitate navigation even for people who would otherwise be bad navigators.

Methods

Participants

Sixty healthy individuals were recruited for the experiment and randomly assigned in equal numbers to three experimental conditions (open maze, closed maze, or teleport maze; 20 participants each). Data from 5 participants were excluded before analysis: 3 because they failed to complete the spatial learning task in the allotted time, 1 because they failed to complete the spatial learning task due to nausea, and 1 because their data were lost due a technical error. Five individuals were recruited to replace these participants and assigned to the corresponding conditions. Of the 60 participants whose data are reported, 39 identified as female, 21 identified as male, and 1 did not disclose their gender. Mean age was 24 (standard deviation 8.4). All participants provided written informed consent in compliance with procedures approved by the University of Pennsylvania Institutional Review Board. One of the included participants did not have data for the JRD and post-experiment questionnaire due to a technical error, and one did not have data for the PTTA task.

Virtual environments

Three virtual environments (open maze, closed maze and teleport maze) were created using Unity 3D software. Each environment contained nine square (20x20 virtual meter) “rooms”. The rooms were connected to each other, such that direct travel was possible between each room and two or three other rooms. The nature of these connections differed between environments, as described below, but they always involved the same pairs of rooms so that the connectivity structure was the same across all three environments (Figure 1A). Each room had a unique floor color (black, white, red, pink, yellow, orange, green, blue, or purple), and contained a treasure chest mounted on a pedestal, with an object inside it (ruby, globe, flashlight, book, key, burger, sword, rose, or bottle). Floor colors and objects were randomly assigned to the different rooms for each participant. All objects and environmental elements were purchased from Unity Asset Store.

In the open maze, the “rooms” were brick patios without walls. They were laid out on a grassy lawn and were connected by stone pathways (also without walls). Travel was restricted to these pathways. Due to the absence of walls, open visibility was maintained across the entire environment. The lawn was bounded by a low fence that marked off a rectangular area of 130 x 100 virtual meters. Distal landmarks were located beyond the four sides of the rectangle: a rocky mountain range, a storefront, a Ferris wheel and a forest (Figure 1A).

In the closed maze, the rooms were enclosed spaces with bounding walls but no ceiling. They were connected by passageways that also had bounding walls and were open to the sky. There were doors between the rooms and the pathways, so that visibility was always limited to the current room or pathway. There were no distal landmarks (Figure 1B).

In the teleport maze, the rooms were the same as in the closed maze. However, in this case, the rooms had no visible exits. They were connected instead via a system of “teleporters” that were engaged when the participants navigated to an upright 3 virtual meter tall white capsule in the center of each room. Upon reaching the capsule, the names of two or three adjacent rooms, along with corresponding color patches, would appear on the screen. A key press of “1”, “2”, or “3” initiated “teleportation” to the chosen room. During teleportation, a black starry space image appeared on the screen for 2.5 seconds, after which participants landed in a random corner of the target room, facing the center of the room. Thus, teleportation moved participants from one room to another without allowing them to maintain a consistent sense of direction. The connections between rooms through the teleporters were exactly the same inter-room connections as in the open and closed maze environments. As in the closed maze environment, there were no distal landmarks (Figure 1C).

Experimental Procedure Overview

At the beginning of the experiment, participants completed the *Santa Barbara Sense of Direction (SBSOD)* scale and the *Perspective Taking Task for Adults (PTTA)*. Participants assigned to the teleport maze condition then completed a brief slideshow tutorial explaining the teleporter system functionality. All participants then completed the *environmental learning task,* which was intended to teach them about the spatial arrangement of the environment and the location of the objects within it. The knowledge that they obtained was then assessed by a series of spatial memory tasks: *Euclidean distance estimation*, *shortest path selection*, *path distance estimation*, *free recall*, *map localization,* and *Judgment of Relative Direction (JRD)*. Finally, participants filled out a *post-experiment questionnaire* (see Figure S1).

The experiment was performed in a single behavioral session, which was administered virtually while being monitored by the experimenter through the Zoom video conferencing software. The session lasted between 1.5-2 hours (mean = 79.6 minutes, SD = 23.2 minutes). Below we describe each experimental task in detail. All tasks were programmed using Unity 3D software and were self-paced.

Santa Barbara Sense of Direction (SBSOD)

Participants completed the Santa Barbara Sense of Direction (SBSOD) scale, a 15 question self-assessment of their navigational ability (Hegarty, 2002). On each trial, a phrase appeared on the screen, and participants asked to rank how much this phrase describes them on a scale of 1 to 7.

Perspective Taking Task for Adults (PTTA)

Participants completed the Perspective Taking Task for Adults (PTTA), which is designed to evaluate the ability to take a spatial perspective other than one's own (Frick et al., 2014). On each trial, participants viewed an image of a person taking a picture of 3 colored geometric shapes, and were instructed to select one out of eight photos that accurately displays the view from the photographer's perspective. Before the task, participants viewed a brief tutorial slideshow to become familiar with the task. Participants were instructed to complete as many trials as possible within 3 minutes.

Environmental learning task

During the environmental learning task, participants freely navigated the virtual environment, while viewing it from a first-person, ground-level perspective. The task was divided into six stages. Stages 1-3 focused on learning the room locations; stages 4-6 focused on learning the object locations.

In each stage, participants were required to navigate to 9 navigational goals (rooms or objects) in a random sequence. Each stage began with short instructions, after which the name and image of the first navigational goal (room or object) was shown at the top of the screen. Participants were required to navigate to the goal and press the spacebar key when they arrived at its location (inside the goal room, or next to the treasure chest containing the goal object). If they pressed the button at the correct goal location, the next navigational goal was displayed. If they pressed the button at an incorrect location, an error message was displayed, and they had to continue searching until they reached the correct location, at which point they moved on to the next trial. After all nine rooms or objects were found, participants were cued to search again for any items they made errors on, and they only moved on to the next stage after they found all nine goals in errorless trials. A counter at the top of the screen indicated how many rooms or objects had been found successfully during the current stage. Participants were randomly placed in the center of one of nine rooms at the start of stage 1, and started each subsequent stage from the ending position of the previous stage.

In stages 1-3, participants searched for the rooms denoted by their floor colors. In stage 1, each room's floor color was made visible as soon as the participant was within the room boundaries. This meant that they could always see the floor color of the room they were in, but could not see the floor colors of distant rooms in the open maze, even though the rooms themselves were visible. In stage 2, each room's floor color was only visible once the participants indicated that they were at the goal room, forcing participants to use their memory for the floor colors. Stage 3 was similar to stage 2, but all rooms had to be found in errorless sequence in order to finish the task; if an error was made, the stage started over at the beginning, with a new random sequence of trials. Throughout stages 1-3, each treasure chest was opened as soon as participants entered its corresponding room, making the objects visible (even though they were not yet relevant to the task).

In stages 4-6, participants searched for the objects. In stage 4, the treasure chest in each room opened as soon at the participant entered the room, revealing the object within. In stage 5, objects remained hidden inside the treasure chests until participants approached one of the chests and indicated that it was the goal location, at which point the chest opened. Stage 6 was similar to stage 5, but participants were required to find all objects in an errorless sequence, and had to repeat the whole stage from the beginning if they made a mistake. Throughout stages 4-6, floor colors were made visible as soon as participants entered each room.

The learning task was limited to 65 minutes; participants who exceeded this limit were not tested further. The gradual learning, repetition of incorrectly remembered rooms and objects at the end of each stage, and requirement for perfect color/object finding in stages 3 and 6 ensured that participants who completed the learning task accurately encoded all of the room and object locations.

Euclidean distance estimation task

On each trial, the names of two objects were presented, one on the left and one on the right side of the screen. Participants then typed in their estimate of the direct-line (Euclidean) distance between the two objects, in feet. All possible pairs of objects were used, resulting in 36 trials.

Shortest path selection task

On each trial, participants saw the name of a starting room (indicated by floor color) and the name of a target object. Below these on the screen, they saw the names and color patches corresponding to the rooms (two or three) that were connected to the starting room. Their instructions were to choose the connecting room that would take them from the starting room to the target object using the shortest possible path. All possible room-object combinations were queried, with the exception of combinations for which target objects were in rooms adjacent to the starting room and combinations for which there was no correct answer because all selections had a similar shortest path distance. With these exclusions, there were 36 trials.

Path distance estimation task

On each trial, participants viewed the names of two objects presented on the left and right side of the screen. They were instructed to type their estimate of the time (in seconds) it would take them to travel between the two objects. All possible pairs of objects were used, resulting in 36 trials.

Free recall task

Participants were asked to type the names of the objects in the maze in any order, pressing the “return” button after each name to move on to the next line. They then pressed the “finish” button when they had recalled as many objects as possible. Entered object names remained visible along with a counter indicating the number of entered objects.

Map localization task

On each trial, participants were presented with the name and picture of an object on the screen, or the name and picture of a floor color corresponding to one of the rooms, along with an empty rectangle proportional to the environment size. They were instructed to click the cursor within the rectangle to indicate the location of the indicated item, at which point a red dot appeared in the clicked location; participants could click again to reselect the location as many times as they wanted before finalizing their answer by clicking a “continue” button. Each room and object (18 total) was queried, in random order.

Judgment of Relative Direction (JRD) task

On each trial, participants saw the names of two floor colors (corresponding to two rooms) and one object. They were instructed to imagine that they were standing in the first room (starting room), looking toward the second room (facing room). They were then asked to indicate the direction of the object (target object) by rotating an arrow on the screen from 0 to 360 degrees. Each possible starting-facing room combination was queried, for a total of 72 trials. Target objects were never in the starting or facing rooms, and each target object was used an equal number of times. Only participants in the open and closed maze environments completed the JRD task.

QUANTIFICATION AND STATISTICAL ANALYSIS

Santa Barbara Sense of Direction (SBSOD)

Participants' self-ratings in the SBSOD questionnaire were averaged across questions (taking into account questions that are reverse scored). Scores across participants were then correlated to individual performance in each spatial memory task.

Perspective Taking Task for Adults (PTTA)

Scores corresponded to the number of trials answered correctly within the time limit, out of a maximum of 32. Scores across participants were then correlated to performance in each spatial memory task.

Environmental learning task

We calculated a measure of navigational efficiency for each participant in the following manner. First we calculated, for each trial, the length of the path that the participant took from the starting location (i.e. the location of the room/object that was the goal on the previous trial) to the goal location. In the open and closed mazes, the actual path length between the room centers was used (i.e. virtual meters); for the teleport maze condition, the path length was taken to mean the number of rooms through which the participant passed (since all inter-room transitions were of similar length in this condition), and 1 was subtracted from this number (since the floor color of the target room, or the room containing the target object, was always visible upon reaching the teleporter in the room preceding the target). Then, for each trial, chance level performance was calculated by generating 1000 random walks between the starting and goal location and averaging the path lengths of these random walks. The true path length was then divided by the average chance path length to obtain a path efficiency ratio. These ratios were compared to 1 (the null hypothesis of no difference between chance and actual performance) using a one-sample one-tailed t-test across participants in each condition, with FDR-correction across conditions. Stages 1 and 4 were not included in this calculation because paths in these stages were implemented prior to learning room and object locations.

To test for possible use of cognitive graph knowledge, we examined trials for which there were two pathways of equal length but different number of intervening rooms between the starting and target locations. We calculated the proportion of these trials on which participants chose the route with the lower compared to the higher number of intervening rooms. This value was compared to the null hypothesis of no route preference (0.5) using a one-sample two-tailed t-test in each environmental condition, with FDR-correction across conditions. These values were computed for the open and closed maze conditions. They were not calculated in the teleport maze, because path length and number of intervening rooms were identical quantities in this condition.

Euclidean distance estimation task

Task accuracy was computed for each participant by taking the correlation between the estimated distance on each trial and the veridical Euclidean distance. Accuracy was compared to chance for each environment by using a one-sample one-tailed t-test against a zero baseline, FDR-corrected across the three environments. Differences between conditions were tested using a one-way ANOVA with Tukey-Cramer post-hoc tests.

Use of cognitive graph knowledge was evaluated in the following manner. For each participant, we calculated the correlation between the estimated distances across trials and the shortest path length. We compared this value to the correlation between estimated distances and Euclidean distances (i.e. task accuracy) using a paired-samples two-tailed t-test. We reasoned that participants who used graph-knowledge to estimate distances would show greater correlation with path distance than with Euclidean distance. This analysis was not performed for participants in the teleport maze, because these participants only had direct knowledge about the number of intervening rooms, not path length or Euclidean distance.

Shortest path selection task

We calculated the percentage of correct responses for each participant. In the open and closed mazes, correct responses were based on the path length. In the teleport maze, correct responses were based on the number of rooms connecting the starting and target objects. Chance performance was estimated for each maze by making 1000 random answer selections for each question (average chance level accuracy - 0.49 for the open and closed maze conditions, 0.45 for the teleport maze condition). Accuracy was compared to chance using a one-sample one-tailed t-test across participants, with FDR-correction across conditions. Differences between conditions were tested using a one-way ANOVA with Tukey-Cramer post-hoc tests.

To test for use of cognitive graph knowledge, we examined trials in the open and closed mazes for which there were two pathways of equal length but different number of intervening rooms between the starting and target rooms. We calculated the fraction of these trials on which participants chose the route with the lower compared to the higher number of intervening rooms. This value was compared to the null hypothesis of no route preference (0.5) using a one-sample two-tailed t-test, with FDR-correction across mazes.

Path distance estimation task

Task accuracy was computed for each participant by taking the correlation between the estimated distance on each trial and the veridical shortest path distance between objects (for the open and closed maze conditions' participants) or the veridical minimal number of rooms connecting the starting and target objects (for the teleport maze condition participants, since inter-room transitions did not differ in length in this condition). Above-chance performance in each environmental condition was measured using a one-sample one-tailed t-test across participants in each condition against a zero baseline, with FDR-correction across conditions. Differences between conditions were tested using a one-way ANOVA with Tukey-Cramer post-hoc tests.

Free recall task

The relation between recall order and objects' spatial distance was measured by calculating the distance between each pair of consecutively-recalled objects, averaging these values for each participant, and comparing these values to chance-level distances (estimated using 1000 random permutations of object names) using a one-sample one-tailed t-test across participants in each condition, with FDR-correction across conditions. Differences between conditions were tested using a one-way ANOVA with Tukey-Cramer post-hoc tests.

Map localization task

Task accuracy was assessed by taking each participant's localization responses and using a gradient descent algorithm to translate, scale, and rotate these responses by multiples of 90 degrees until they best fit the true configuration. Accuracy was then determined by measuring the average distance between each object's marked location and its true location. These values were compared to chance-level performance (estimated using the same process for 1000 random localizations of 9 items), using a one-sample one-tailed t-test across participants in each condition, with FDR-correction across conditions. Differences between conditions were tested using a one-way ANOVA with Tukey-Cramer post-hoc tests.

Judgment of Relative Direction (JRD) task

Task accuracy was computed as the mean angular distance between participants' responses and the veridical object directions. These values were compared to chance-level performance (90 degrees average deviation) using a one-sample one-tailed t-test across participants in each condition, with FDR-correction across conditions. The difference between the open and closed maze conditions was tested using a two-sample two-tailed t-test.

Individual variability analysis

To investigate individual variability in environmental learning, we used the map localization task results to assign participants to groups, because this task provides the most direct test of participants' knowledge of environmental layout. We first divided the 60 participants into two groups based on the median value (1.29) of their performance on map localization task. To ensure that these groups reflected two distinct clusters, the silhouette value of the grouping was computed, and this value was compared to the silhouette value of randomly-generated data in the same value range that were divided to two groups using the median, repeated 1000 times. The silhouette of the original grouping was higher than all random data grouping, indicating separation between groups at $p<0.001$. Further validating this grouping, K-means clustering (K=2) of the data led to identical assignment of participants into two clusters. The two groups of participants (30 participants in each) were subsequently named "integrators" and "non-Integrators" (following Weisberg & Newcombe 2018). Within the integrators group, 18 participants were in the open maze, 9 participants were in the closed maze, and 3 participants in the teleport maze. Within the non-Integrators group, 2 participants were in the open maze, 11 participants were in the closed maze, and 17 participants were in the teleport maze. To assess the difference in distribution of integrators and non-integrators between the groups, we used a Chi-square test.

To investigate the effect of individual variability on task performance, we analyzed the accuracy of integrators and non-integrators in the Euclidean distance estimation task, shortest path selection task, path distance estimation task, and JRD, in each environment separately. Since there were few non-integrators in the open maze or integrators in the teleport maze, these two groups (open maze non-integrators and teleport maze integrators) were omitted from the analyses. Differences between the remaining four participant groups (open maze integrators, closed maze integrators, closed maze non-integrators and teleport maze non-integrators) were tested for each task using a one-way ANOVA with Tukey-Cramer post-hoc tests. The resulting p-values were FDR-corrected across tasks.

DATA AVAILABILITY

Data and analysis codes for this project are publicly available at: https://osf.io/adcfk/?view_only=28a03c1cd2314bd6995a7c5ca2e69180.

Acknowledgments

M.P. is supported by the Fulbright and Zuckerman STEM Leadership fellowships. This research was supported by NIH grants R01EY022350 and R01EY031286 to R.A.E. We thank Alexander Miller and the UPenn VR lab for providing tools for Unity data logging.

Author contributions

M.P., C.N. and R.A.E – conceptualization, data curation, formal analysis, investigation, methodology, validation, visualization, writing (original draft, review and editing). M.P. and C.N. – software. R.A.E – funding acquisition, project administration and supervision.

References

Allen, G. L., Kirasic, K. C., Dobson, S. H., Long, R. G., & Beck, S. (1996). Predicting environmental learning from spatial abilities: An indirect route. *Intelligence*, *22*(3), 327–355. https://doi.org/10.1016/S0160-2896(96)90026-4

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*, 101580. https://doi.org/10.1016/j.jenvp.2021.101580

Chadwick, M. J., Jolly, A. E. J., Amos, D. P., Hassabis, D., & Spiers, H. J. (2015). A Goal Direction Signal in the Human Entorhinal/Subicular Region. *Current Biology*, *25*(1), 87–92. https://doi.org/10.1016/j.cub.2014.11.001

Chrastil, E. R., & Warren, W. H. (2014). From cognitive maps to cognitive graphs. *PloS One*, *9*(11), e112544.

Coutrot, A., Manley, E., Goodroe, S., Gahnstrom, C., Filomena, G., Yesiltepe, D., Dalton, R. C., Wiener, J. M., Hölscher, C., Hornberger, M., & Spiers, H. J. (2022). Entropy of city street networks linked to future spatial navigation ability. *Nature*, *604*(7904), 104–110. https://doi.org/10.1038/s41586-022-04486-7

Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, *463*(7281), 657–661. https://doi.org/10.1038/nature08704

Doner, S., Zheng, J., McAvan, A. S., Starrett, M. J., Campbell, H., Sanders, D., & Ekstrom, A. (2022). Evidence for flexible navigation strategies during spatial learning involving path choices. *Spatial Cognition & Computation*, 1–30. https://doi.org/10.1080/13875868.2022.2158090

Ericson, J. D., & Warren, W. H. (2020). Probing the invariant structure of spatial knowledge: Support for the cognitive graph hypothesis. *Cognition*, *200*, 104276. https://doi.org/10.1016/j.cognition.2020.104276

Etienne, A. S., & Jeffery, K. J. (2004). Path integration in mammals. *Hippocampus*, *14*(2), 180–192. https://doi.org/10.1002/hipo.10173

Fields, A. W., & Shelton, A. L. (2006). Individual skill differences and large-scale environmental learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(3), 506–515. https://doi.org/10.1037/0278-7393.32.3.506

Frick, A., MÃ¶hring, W., & Newcombe, N. S. (2014). Picturing perspectives: Development of perspective-taking abilities in 4- to 8-year-olds. *Frontiers in Psychology*, *5*. https://doi.org/10.3389/fpsyg.2014.00386

Gallistel, C. R. (1990). *The organization of learning* (pp. viii, 648). The MIT Press.

He, Q., & Brown, T. I. (2019). Environmental Barriers Disrupt Grid-like Representations in Humans during Navigation. *Current Biology*, *29*(16), 2718-2722.e3. https://doi.org/10.1016/j.cub.2019.06.072

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

Hirtle, S. C., & Jonides, J. (1985). Evidence of hierarchies in cognitive maps. *Memory & Cognition*, *13*(3), 208–217. https://doi.org/10.3758/BF03197683

Ishikawa, T., & Montello, D. (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

Jacobs, J., Weidemann, C. T., Miller, J. F., Solway, A., Burke, J. F., Wei, X.-X., Suthana, N., Sperling, M. R., Sharan, A. D., Fried, I., & Kahana, M. J. (2013). Direct recordings of grid-like neuronal activity in human spatial navigation. *Nature Neuroscience*, *16*(9), 1188–1190. https://doi.org/10.1038/nn.3466

Kozhevnikov, M., Motes, M. A., Rasch, B., & Blajenkova, O. (2006). Perspective-taking vs. Mental rotation transformations and how they predict spatial navigation performance. *Applied Cognitive Psychology*, *20*(3), 397–417. https://doi.org/10.1002/acp.1192

Kuipers, B. (1982). The “Map in the Head” Metaphor. *Environment and Behavior*, *14*(2), 202–220. https://doi.org/10.1177/0013916584142005

Maidenbaum, S., Miller, J., Stein, J. M., & Jacobs, J. (2018). Grid-like hexadirectional modulation of human entorhinal theta oscillations. *Proceedings of the National Academy of Sciences*, *115*(42), 10798–10803. https://doi.org/10.1073/pnas.1805007115

McNaughton, B. L., Barnes, C. A., Gerrard, J. L., Gothard, K., Jung, M. W., Knierim, J. J., Kudrimoti, H., Qin, Y., Skaggs, W. E., Suster, M., & Weaver, K. L. (1996). Deciphering the hippocampal polyglot: The hippocampus as a path integration system. *Journal of Experimental Biology*, *199*(1), 173–185. https://doi.org/10.1242/jeb.199.1.173

McNaughton, B. L., Barnes, C. A., & 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*(1). https://doi.org/10.1007/BF00237147

McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M.-B. (2006). Path integration and the neural basis of the "cognitive map." *Nature Reviews Neuroscience*, *7*(8), Article 8. https://doi.org/10.1038/nrn1932

Mehta, M. R., Barnes, C. A., & McNaughton, B. L. (1997). Experience-dependent, asymmetric expansion of hippocampal place fields. *Proceedings of the National Academy of Sciences*, *94*(16), 8918–8921. https://doi.org/10.1073/pnas.94.16.8918

Meilinger, T. (2008). The Network of Reference Frames Theory: A Synthesis of Graphs and Cognitive Maps. In C. Freksa, N. S. Newcombe, P. Gärdenfors, & S. Wölfl (Eds.), *Spatial Cognition VI. Learning, Reasoning, and Talking about Space* (pp. 344–360). Springer. https://doi.org/10.1007/978-3-540-87601-4_25

Miller, J. F., Lazarus, E. M., Polyn, S. M., & Kahana, M. J. (2013). Spatial clustering during memory search. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(3), 773–781. https://doi.org/10.1037/a0029684

Moeser, S. D. (1988). Cognitive Mapping in a Complex Building. *Environment and Behavior*, *20*(1), 21–49. https://doi.org/10.1177/0013916588201002

Moser, E. I., Kropff, E., & Moser, M.-B. (2008). Place Cells, Grid Cells, and the Brain's Spatial Representation System. *Annual Review of Neuroscience*, *31*(1), 69–89. https://doi.org/10.1146/annurev.neuro.31.061307.090723

Muller, R., Bostock, E., Taube, J., & Kubie, J. (1994). On the directional firing properties of hippocampal place cells. *The Journal of Neuroscience*, *14*(12), 7235–7251. https://doi.org/10.1523/JNEUROSCI.14-12-07235.1994

Muryy, A., & Glennerster, A. (2018). Pointing Errors in Non-metric Virtual Environments. In S. Creem-Regehr, J. Schöning, & A. Klippel (Eds.), *Spatial Cognition XI* (pp. 43–57). Springer International Publishing. https://doi.org/10.1007/978-3-319-96385-3_4

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, *34*(1), 171–175. https://doi.org/10.1016/0006-8993(71)90358-1

O'Keefe, J., & Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press. https://repository.arizona.edu/handle/10150/620894

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

Peer, M., & Epstein, R. A. (2021). The human brain uses spatial schemas to represent segmented environments. *Current Biology*, *31*(21), 4677-4688.e8. https://doi.org/10.1016/j.cub.2021.08.012

Schinazi, V. R., Nardi, D., Newcombe, N. S., Shipley, T. F., & Epstein, R. A. (2013). Hippocampal size predicts rapid learning of a cognitive map in humans. *Hippocampus*, *23*(6), 515–528. https://doi.org/10.1002/hipo.22111

Shine, J. P., Valdés-Herrera, J. P., Tempelmann, C., & Wolbers, T. (2019). Evidence for allocentric boundary and goal direction information in the human entorhinal cortex and subiculum. *Nature Communications*, *10*(1), 4004. https://doi.org/10.1038/s41467-019-11802-9

Siegel, A. W., & White, S. H. (1975). The Development of Spatial Representations of Large-Scale Environments. In H. W. Reese (Ed.), *Advances in Child Development and Behavior* (Vol. 10, pp. 9–55). JAI. https://doi.org/10.1016/S0065-2407(08)60007-5

Warren, W. H. (2019). Non-Euclidean navigation. *Journal of Experimental Biology*, *222*(Suppl 1). https://doi.org/10.1242/jeb.187971

Weisberg, S. M., & Newcombe, N. S. (2018). Cognitive Maps: Some People Make Them, Some People Struggle. *Current Directions in Psychological Science*, *27*(4), 220–226. https://doi.org/10.1177/0963721417744521

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), 669–682. https://doi.org/10.1037/a0035261

Zetzsche, C., Wolter, J., Galbraith, C., & Schill, K. (2009). Representation of space: Image-like or sensorimotor? *Spatial Vision*, *22*(5), 409–424. https://doi.org/10.1163/156856809789476074

Supplementary Materials

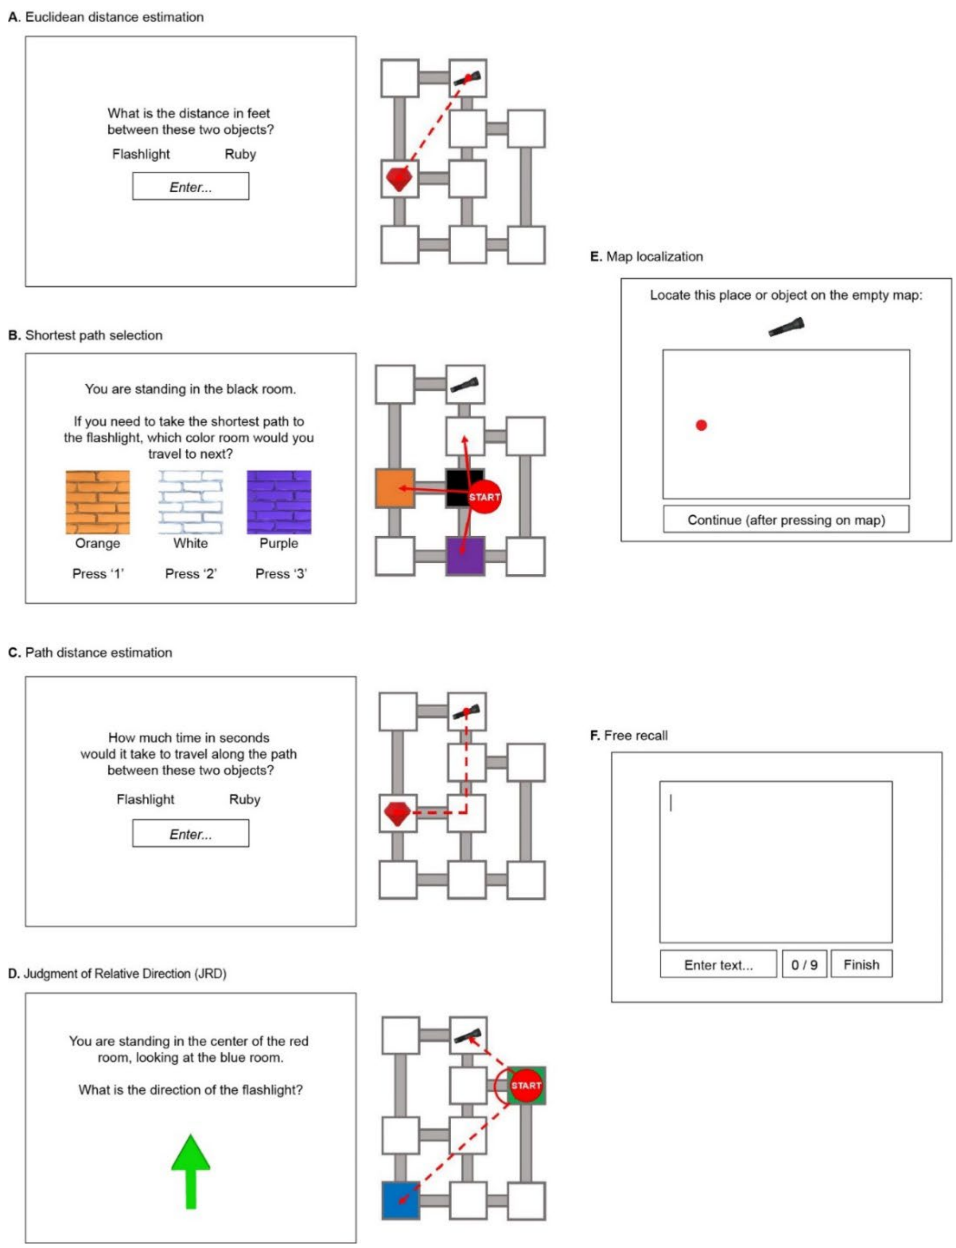

Figure S1: Spatial memory tasks. A-D) Euclidean distance estimation, shortest path selection, path distance estimation, and judgment of relative direction (JRD) tasks. Left – Example trial screen for each task; Right – Schematic depicting the choice the participant needs to make in an example trial. E-F) Map localization and Free recall tasks: Example trial screens.

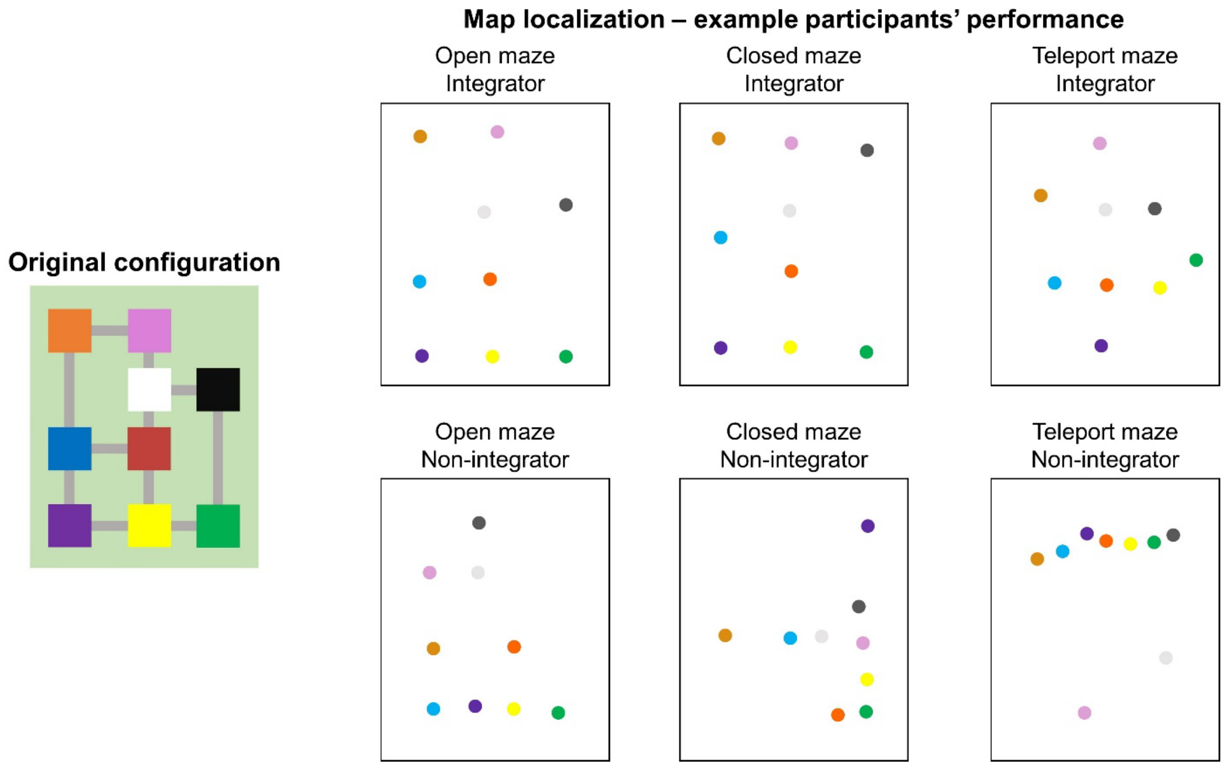

Figure S2: Map localization example performance. Left – the original configuration of rooms in the environment. Right – example configurations rooms in the map localization task, from six example participants (three integrators and three non-integrators, from the three experimental environments; each rectangle represents one participant's performance). Each colored dot represents a participant's localization of the correspondingly colored room on a blank map of the environment. The localization of the nine objects task is not presented here.

	Correlation of individual SBSOD score to task performance					
	Learning	Euclidean distance estimation	Path distance estimation	Shortest path selection	Map localization	JRD
Open maze	0.08	0.14	0.25	0.08	0.11	0.18
Closed maze	0.56	0.34	0.40	0.46	0.44	0.30
Teleport maze	0.21	0.10	0.33	0.52	0.37	
	Correlation of individual PTTA score to task performance					
	Learning	Euclidean distance estimation	Path distance estimation	Shortest path selection	Map localization	JRD
Open maze	0.18	0.25	0.20	0.12	0.67	0.52
Closed maze	0.54	0.45	0.55	0.35	0.46	0.57
Teleport maze	0.18	0.49	0.33	0.51	0.38	

Table S1: Correlation of SBSOD and PTTA scores to task performance.