



Original Articles

Schematic representations of local environmental space guide goal-directed navigation



Steven A. Marchette*, Jack Ryan, Russell A. Epstein

Department of Psychology, University of Pennsylvania, United States

ARTICLE INFO

Article history:

Received 17 March 2016

Revised 27 September 2016

Accepted 7 October 2016

Keywords:

Navigation

Allocentric

Spatial memory

Environmental geometry

Place recognition

Virtual reality

ABSTRACT

To successfully navigate to a target, it is useful to be able to define its location at multiple levels of specificity. For example, the location of a favorite coffee mug can be described in terms of which room it is in, or in terms of where it is within the room. An appealing hypothesis is that these levels of description are retrieved from memory by accessing the same representation at progressively finer levels of granularity—first remembering the general location of an object and then “zooming in.” Here we provide evidence for an alternative view, in which navigational behavior is guided by independent representations at multiple spatial scales. Subjects learned the locations of objects that were positioned within four visually distinct but geometrically similar buildings, which were in turn positioned within a broader virtual park. They were then tested on their knowledge of object location by asking them to navigate to the remembered location of each object. We examined errors during the test phase for confusions among geometrically analogous locations in different buildings—that is, navigating to the right location in the wrong building. We observed that subjects frequently made these confusions, which are analogous to remembering a passage’s location on the page of a book but not remembering the page that the passage is on. This suggests that subjects were recalling the object’s local location *without* recalling its global location. Further manipulations across seven experiments indicated that geometric confusions were observed even between buildings that were not metrically identical as long as geometrical equivalence could be defined. However, removing the walls so that the larger environment was no longer divided into subspaces abolished these errors. Taken together, our results suggest that human spatial memory contains two separable representations of “where” an object can be found: (i) a schematic map of where an object lies with respect to local landmarks and boundaries; (ii) a representation of the identity and location of each local environment.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Space has behavioral relevance at several different scales. For example, the location of a coffee mug sitting on a desk can be described in terms of where it is positioned on the desk, where the desk is in the room, where the room is within the building, and where the building is situated in the wider world (Montello, 1993). To successfully navigate to a target, one must be able to access spatial memories at several of these scales: it does one little good to know where the mug is in one’s office if one cannot find the office or (even worse) cannot remember which building the office is in. The current study attempts to understand how we bridge these spatial scales during goal-directed navigation.

Although human spatial memory has been consistently observed to have a hierarchical organization involving multiple scales (e.g., Hirtle & Jonides, 1985; Huttenlocher, Newcombe, & Sandberg, 1994; Kosslyn, Pick, & Fariello, 1974; McNamara, Hardy, & Hirtle, 1989), to our knowledge, most theories of spatial memory are agnostic on how these multiple scales are accessed during retrieval. An appealing hypothesis is that we accomplish this task by accessing the same representation at progressively finer levels of granularity. For example, when remembering how to get to the coffee cup, we first remember its general location and then “zoom in” to determine its precise coordinates. The result of such a coarse-to-fine strategy would be *context-specific* memory retrieval: an object’s fine-scale spatial coordinates would be recallable only if the correct environmental context (i.e. its location at a coarse scale) were also remembered. Evidence for such a top-down approach to encoding and retrieval of spatial information comes from observations that memory for an object’s precise location is

* Corresponding author at: Department of Psychology, University of Pennsylvania, 3720 Walnut Street, Philadelphia, PA 19104, United States.

E-mail address: stmar@sas.upenn.edu (S.A. Marchette).

sometimes biased towards the center of the local division containing it (Hirtle & Jonides, 1985; Huttenlocher, Hedges, & Duncan, 1991; Huttenlocher et al., 1994; Kosslyn et al., 1974; McNamara et al., 1989; Newcombe, Huttenlocher, Sandberg, Lie, & Johnson, 1999; Stevens & Coupe, 1978). Although this does not necessitate that retrieval of spatial information on the coarse scale precedes retrieval on the fine scale, it does suggest that knowledge about coarse-scale divisions can guide retrieval of fine-scale spatial information and may logically precede this retrieval (Huttenlocher et al., 1991).

However, even if spatial memory recall proceeds in a coarse-to-fine manner in normal circumstances, this does not necessarily mean that retrieval of spatial information at the coarse level and retrieval at the fine level are merely two temporal stages of the same mental process. It is possible that these two levels of spatial recall may be distinct mental operations that can be decoupled from each other (Meilinger, 2008). A common experience that illustrates this phenomenon is remembering a passage's location on the page of a book without remembering the page itself. In the navigational realm, the equivalent error would be remembering where an object is within a room without remembering the room containing it. Both are examples of remembering where an object is within a context without remembering which context it is in.

There is some neuroscientific data that suggests that this kind of decoupling may be possible, though the results are mixed, with some results more consistent with the first (coarse-to-fine) view, and some more consistent with the second view (distinct mental operations). Supporting the first view, hippocampal place cells, which are believed to neurally instantiate memories for spatial locations (Miller et al., 2013), often code space in a context-specific manner: place cells fire at a specific location within an environment (in rats: O'Keefe & Dostrovsky, 1971; and in humans: Ekstrom et al., 2003), but *remap* to have different firing fields in different environments (Colgin, Moser, & Moser, 2008; Jeffery & Anderson, 2003; Muller & Kubie, 1987). This remapping phenomenon is consistent with a coarse-to-fine scheme, insofar as it implies that identifying the current environmental context—and retrieval of the unique “map” that is specific to that context—is a necessary precursor to recall of a precise spatial location. Supporting the alternative (distinct mental operations) view, place cells have also been observed to fire in geometrically equivalent locations when animals explore different environments (Lever, Wills, Cacucci, Burgess, & O'Keefe, 2002; O'Keefe & Burgess, 1996; Wills, Lever, Cacucci, Burgess, & O'Keefe, 2005), suggesting that there are some circumstances under which spatial locations can be represented in a manner that is *not* unique to the current local context. Recent evidence also suggests that these geometric repetitions may also occur when animals travel directly from one part of an environment to another. For example, Spiers, Hayman, Jovalekic, Marozzi, and Jeffery (2015) found that when rats explored an environment consisting of several connected chambers of the same size and shape, hippocampal firing fields did not remap across chambers, but rather repeated across chambers in geometrically equivalent positions. Similarly, in a recent fMRI study from our laboratory, we used multivoxel pattern analysis to show that the retrosplenial complex (a brain region that—like the hippocampus—is believed to be crucial for spatial memory) exhibited generalization across geometrically equivalent locations in different virtual “museums” when subjects recalled spatial memories (Marchette, Vass, Ryan, & Epstein, 2014). In other words, this brain region coded location within the local environment without distinguishing between different local environments that had the same size and shape. Taken together, these results provide some evidence for coding of local spatial coordinates in a manner

that is at least partially independent of the code for environmental identity.

Although suggestive, these neuroscientific data are not strong enough to be adjudicative. What is missing from the literature is direct behavioral evidence for memory *confusions* between analogous locations in different environments. This error should be possible if object locations are encoded using a context-independent reference frame that represents only the local spatial structure, because locations in different environments should receive the same mental coordinates if they are located in analogous positions with respect to the surrounding local landmarks and geometry. We would conceptualize such a representation as analogous to a schematic drawing—or blueprint—of the local environment because it preserves the spatial relationships among its essential parts without explicitly representing the specific details that distinguish the environment from other spatially similar environments. We refer to such a representation as a *spatial schema* because it extracts the essential spatial elements of the environment (Bartlett, 1932; Piaget, 1929; Tse et al., 2007; van Buuren et al., 2014) and can potentially be applied to multiple environments. Memory confusions between analogous locations should be unlikely, on the other hand, if locations are represented with respect to context-specific reference frames because in this case an object's location could not be specified without first identifying its environmental context.

During the training phase of our aforementioned fMRI study (Marchette et al., 2014), we observed anecdotal evidence for exactly this kind of confusion. The virtual environment that subjects learned consisted of four buildings arranged in a courtyard, each of which had a single geometrically equivalent interior room that contained objects to be remembered located along the walls. After a brief period of free exploration, subjects were returned to the starting point and instructed to navigate to the remembered location of each object. Although they were usually accurate at this task, on some trials we observed an intriguing error: participants would navigate to the right location in the wrong building before realizing that they had arrived at the wrong object and correcting their mistake. This suggests that their initial route was guided by correct knowledge of where the target object was located relative to the local reference frame, but incorrect knowledge about the identity of the local environment.

The current experiment attempts to systematize these observations of confusions among geometrically analogous locations during navigation, and by doing so test the existence and characteristics of context-independent spatial reference frames. To do this, we once again taught participants a virtual environment that contained four geometrically equivalent buildings (Marchette et al., 2014). We then tested their ability to navigate to the remembered locations of the objects after they had been removed. We reasoned that if participants encoded all four buildings using the same schematic reference frame, they would occasionally confuse a target's location with analogous locations in the other buildings. Such errors would be evidence that subjects had successfully retrieved the location of the object relative to the local reference frame without retrieving the identity of the local environment—thus indicating a potential dissociation between these two kinds of spatial codes.

2. Method

2.1. Participants

168 participants (84 female) from the University of Pennsylvania participated in exchange for course credit or monetary reward. All participants provided written informed consent in compliance with procedures approved by the University of Pennsylvania

Institutional Review Board. We performed seven experiments, each involving 24 participants. Gender was equally balanced in all experiments.

2.2. Virtual environments

Participants learned one of seven virtual environments that were modifications of the environment reported in a previous study (Marchette et al., 2014). All environments were constructed with the Source SDK Hammer Editor (www.valvesoftware.com, Valve Software, Bellevue, WA) and rendered and displayed on a laptop running the commercial video game software Portal (www.valvesoftware.com, Valve Software, Bellevue, WA). The basic version of this environment consisted of a park containing four large rectangular buildings, termed “museums,” that were visually distinguished by unique textures and architectural features (e.g. facades and columns) and were geometrically equivalent (Fig. 1). Each museum contained eight distinct, nameable objects (e.g., “Cake”, “Computer”, “Bicycle”, “Car”, etc.) that were arranged in alcoves along the walls (Fig. 1B). Museums were placed in the environment in a clover-leaf pattern so that each one was entered from a unique direction, so that local spatial codes (‘back-left corner’) were dissociated from global spatial codes (‘north-west corner’). Distal landmarks surrounded the park, including a mountain range and the sun at a high azimuth to the “North”, apartment buildings and refineries to the East and West, and two high rise apartment buildings to the South.

This basic version of the virtual environment was used in Experiment 1 to test whether participants made schema-preserving errors (that is, remembered the objects as being in the correct location but in the wrong museum). Modified versions were used in Experiments 2–7 (Figs. 2 and 4) to test alternative hypotheses about the processes that induced these errors and to understand the nature of the schematic representation. Experiment 2 tested whether participants made schema-preserving errors because they remembered a fixed path from the museum door to the target object. To rule out this possibility, participants were required to enter the museum through different doors during the free exploration and navigational testing phases. Experiments 3–5 tested the role of environmental geometry in creating schematic representations; specifically, these experiments tested whether schema-preserving errors only occurred between geometrically identical museums or were possible for geometrically non-identical museums as long as geometrical equivalence could be defined. To this end, we manipulated the geometry of the museums so that half were one shape and the other half were a different shape (Fig. 2B–D); we then observed whether subjects confused locations across museums that were round vs. rectangular (Experiment 3), square vs. rectangular (Experiment 4), or narrow rectangles vs. wide rectangles (Experiment 5). Experiments 6 tested whether bounding geometry was necessary for participants to make schema-preserving errors. The walls and internal features were removed from all museums but the alcoves containing the objects were left in place (Fig. 4A). Finally, Experiment 7 tested whether non-geometric cues that anchor a reference frame could functionally replace the bounding walls. In this experiment, the open-field environment of Experiment 6 was used but large landmarks (e.g. a watertower) were added to the center of each group of alcoves (Fig. 4B). Besides these modifications, the environments in Experiments 2–6 were identical to that used in Experiment 1.

2.3. Procedure

2.3.1. Free exploration

Participants began the experiment at the entrance at the base of the park (labelled “start” in Fig. 1) and were instructed on the use

of arrow keys to navigate the virtual environment. They were then given fifteen minutes to freely explore the park and the museums within it. In Experiments 1–5 participants were instructed to visit each of the museums and view all of the objects contained therein during this time; in Experiments 6 and 7 which used open-field environments in which there were no “museums”, participants were instructed to view all of the individual objects. All participants complied with these instructions. After fifteen minutes, participants were teleported to the park’s entrance and the navigation test began.

2.3.2. Navigation test

During the navigation test, all objects were removed from their alcoves and participants were asked to sequentially locate the objects from memory. At the beginning of each trial, text appeared on the screen prompting the participant to navigate to a specific target object (e.g. “Please find the Cake”). To ensure that participants did not forget the identity of the current target, the name of the object was displayed on the screen until the end of the trial. Participants were instructed to navigate as directly as they could to the alcove containing the prompted object. Upon reaching an alcove, the participants’ response was recorded and they were teleported to the entrance of the park and prompted to find another target object. No time limit was imposed, and participants were free to navigate to the target as they saw fit. Specifically, participants were able to enter and leave museums at will so that they had the opportunity to determine whether they were in the correct museum based upon its interior appearance and adjust their course if they believed they were in the process of making a mistake. No feedback was given as to whether participants had navigated to a correct or incorrect location. Each object was queried once, in random order, for a total of 32 trials.

2.4. Analysis of navigation errors

To analyze participants’ behavior, we scored each trial on the navigation test by comparing the alcove selected by the participant to true location of the target object. A trial was scored as correct if the participant selected the target’s location and incorrect if a different location was selected. We then classified each incorrect trial according to three potential types of error: (i) context-preserving errors in which the participant navigated to the correct museum but wrong alcove within the museum; (ii) schema-preserving errors in which the participant navigated to the alcove that was “correct” as defined by its local position relative to the walls and main entrance but within the wrong museum; (iii) non-systematic errors in which neither local location nor museum were correct. Because we were interested in the proportions of these different types of errors, we could not analyze participants who located every object correctly and thus made no errors. As a result, a total of 7 participants were excluded with between 0 and 2 participants excluded from each experiment (Experiment 1: 2; Experiment 3: 1; Experiment 4: 2; Experiment 6: 2; none were excluded from Experiments 2, 5 and 7).

For each participant, we calculated their total number of errors and the percentage of these errors classified as context-preserving, schema-preserving or non-systematic. To assess significance, we compared these percentages against the baseline chance of making that kind of error. However, we could not directly compare the error rates for different types of error because they had different baseline probabilities. That is, because each of the 4 museums contained 8 objects, a target object was always in the same museum as 7 other objects (potential context-preserving errors) but only in the same local position as 3 other objects (potential schema-preserving errors); and 21 objects shared neither location nor museum with the target (potential non-systematic errors). Thus,



Fig. 1. Map and images of the virtual environment. (A) Map of the virtual park and the four museums used in Experiment 1. Each museum was oriented at a unique direction with respect to the surrounding park. Objects were displayed within alcoves, which are indicated by grey squares. Each alcove could only be viewed from one direction. The yellow alcove represents an example target object; cyan alcoves represent schematic errors for this target; magenta alcoves represent contextual errors for this target. (B) Images of the four museum exteriors, interiors, and an example object in an alcove.

a navigator who randomly selected an alcove would be more likely to make a context-preserving error than a schema-preserving error. To allow comparisons among different types of error, we divided the percentage of errors of each type by their respective base rates. This produces a measure of the proportion of errors of a particular type relative to what would be expected due to chance. For example, if 20% of a participant's errors were schema-preserving then they would score a 2.07 because they made approximately twice as many errors as would be expected by chance (schema-preserving error base rate: $3/31 = 9.6\%$). We then used the normalized proportions of schematic- and context-preserving errors to test whether participants had a greater tendency to make one error or the other.

One possible objection to using the percentage of errors that were schema- or context-preserving errors is that the measure

may place an undue weight on participants who make very few errors. For example, a participant who makes only a single schema-preserving error would score 100% for schema-preserving errors despite providing less compelling evidence for a context-independent representation of location than a participant who made seven schema-preserving errors out of ten total errors (70%). To ensure that this could not account for our effects, we also analyzed these results using the number of errors observed of each type and compared these against the number that would be expected due to chance, and observed the same results. The results with this measure were consistent with the results with the normalized proportion measure, so for brevity we only report the former.

As an additional aid to visualize the pattern of results, we also created diagrams depicting subject's errors in bounded (Experi-

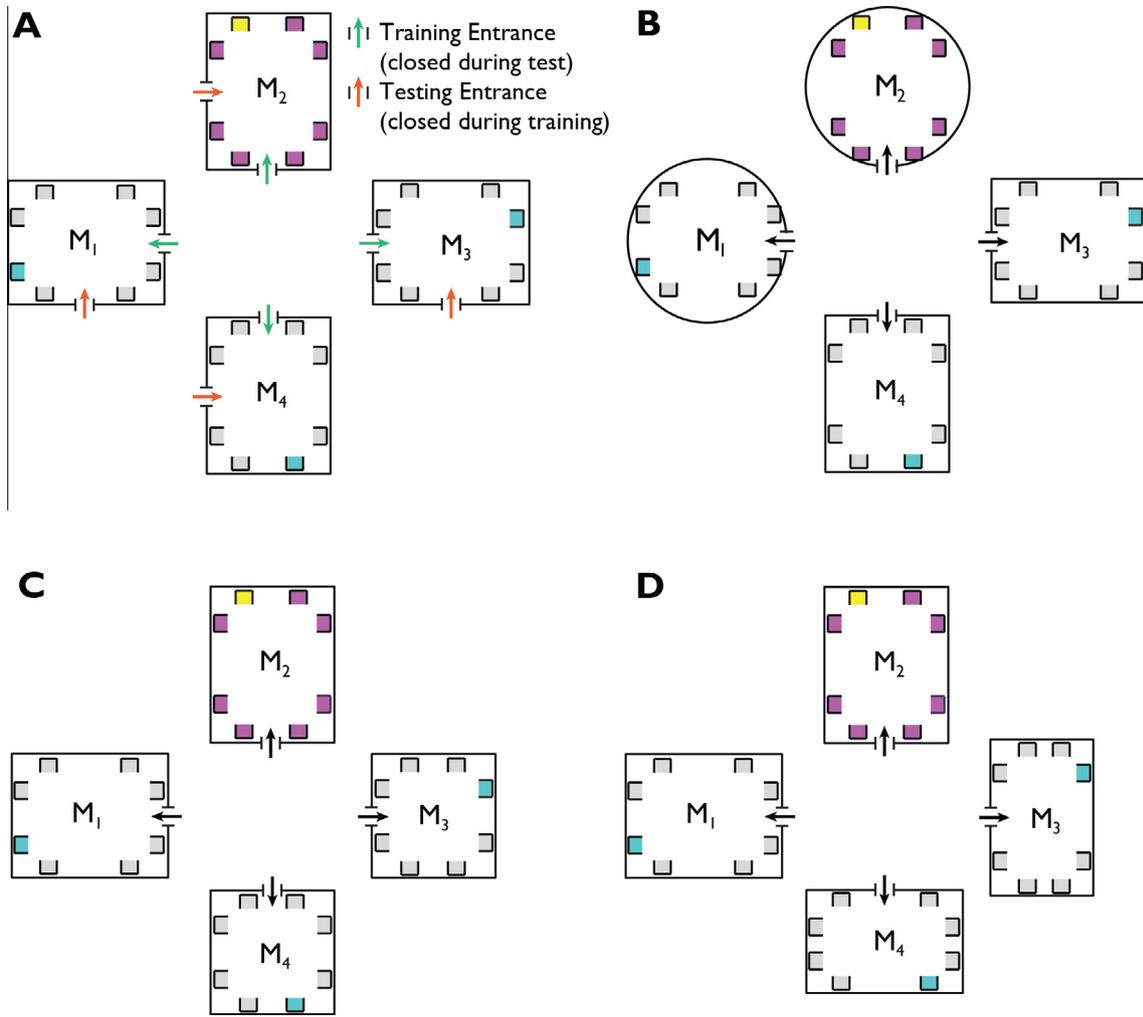


Fig. 2. Illustration of the manipulations of environmental geometry in Experiments 2–5. The yellow alcove represents an example target object; cyan alcoves represent schematic errors for this target; magenta alcoves represent contextual errors for this target. (A) In Experiment 2, participants entered the museum through different doors in the free exploration (green arrow) and navigation testing (orange arrow) phases. (B–D) In Experiments 3–5, the museums were divided into two subsets of museums with different geometries: circles and rectangles in Experiment 3 (B), squares and rectangles in Experiment 4 (C), and narrow and wide rectangles in Experiment 5 (D).

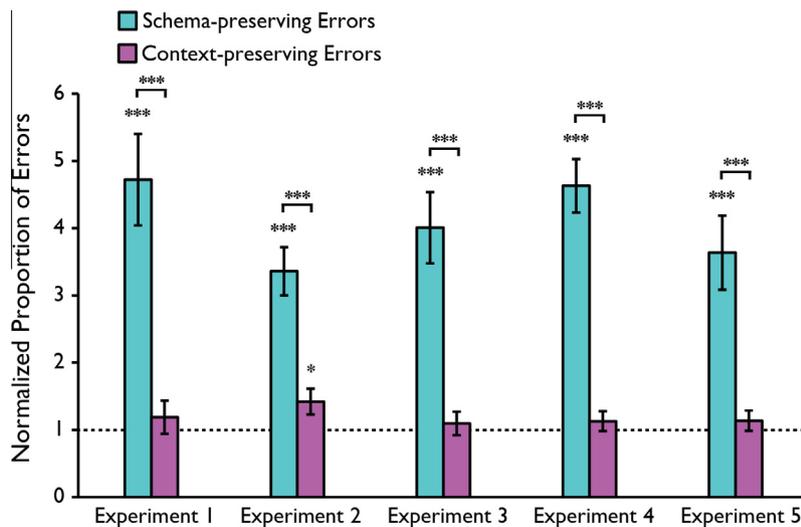


Fig. 3. Proportion of schema-preserving (cyan bars) and context-preserving (magenta bars) errors relative to guessing baseline for Experiments 1–5. Values above 1 indicate an error was made more frequently than expected by chance (dashed line). Error bars represent the standard error of the mean.

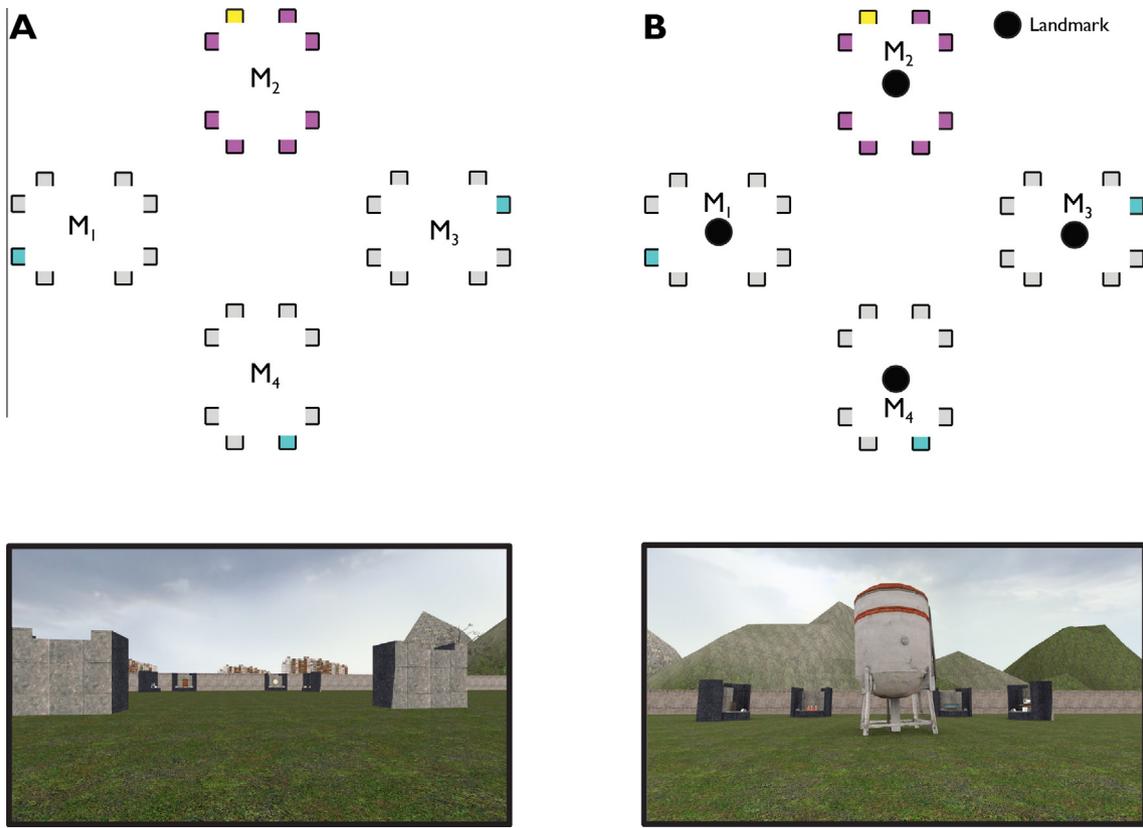


Fig. 4. Map and images of the open-field environments used in Experiments 6 & 7. (A) In Experiment 6, the museum walls and doors were removed but the object alcoves were left in place. (B) In Experiment 7, landmarks were added to the center of each cluster of alcoves.

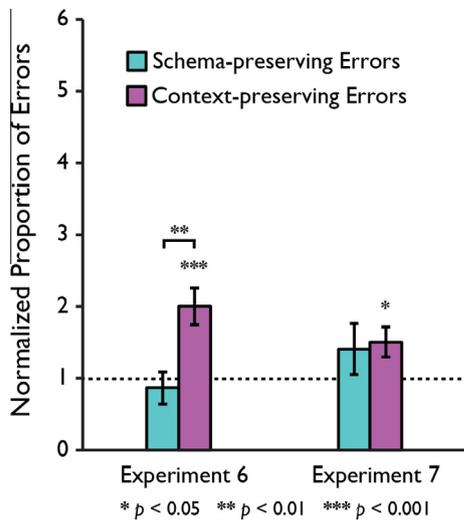


Fig. 5. Proportion of schema-preserving (cyan bars) and context-preserving (magenta bars) errors relative to guessing baseline for Experiments 6 & 7. Values above 1 indicate an error was made more frequently than expected by chance (dashed line). Error bars represent the standard error of the mean.

ments 1–5; Fig. 6) and open-field (Experiments 6 & 7; Fig. 7) environments. To create these diagrams, we first considered each of the 32 target objects in sequence. For each target object, we counted how frequently subjects confused the target object’s location with each of the other locations in the environment. We then used this data to create a map of all of the errors that participants made while attempting to navigate to that target object. Once we had created the 32 error maps (one for each target object), the error

maps were aligned to one another using transformations that maintained the relative spatial relationships between the correct target location and all incorrect locations. The aligned error maps were then summed together and normalized by the total number of errors to create the final diagram. This diagram reflects the percentage of errors that the subjects made at each location, but realigned so that correct responses were all plotted in the same location.

3. Results

3.1. Experiment 1: Schematic representations of spatial location

In the first experiment (Experiment 1), we tested whether participants form schematic representation of object locations while learning a virtual environment consisting of four “museums” of identical geometry. We reasoned that systematic errors during recall could serve as a window into how an object’s location was represented. In particular, two potential types of error could reveal preserved memory for aspects of an object’s location even on incorrect trials. First, *context-preserving errors*—in which subjects navigate to an incorrect location within the correct museum—would indicate preserved memory for the identity of the environmental context containing the target (Fig. 1A). Second, *schema-preserving errors*—in which subjects navigate to the correct location in the wrong museum—would indicate preserved memory for an object’s spatial location within the local context. Finally, subjects could make *non-systematic errors* by navigating to a location that was neither in the correct museum nor in the correct local location.

We first calculated the number of goals correctly located to determine whether participants had successfully encoded the environment during their 15 min of free exploration. Participants

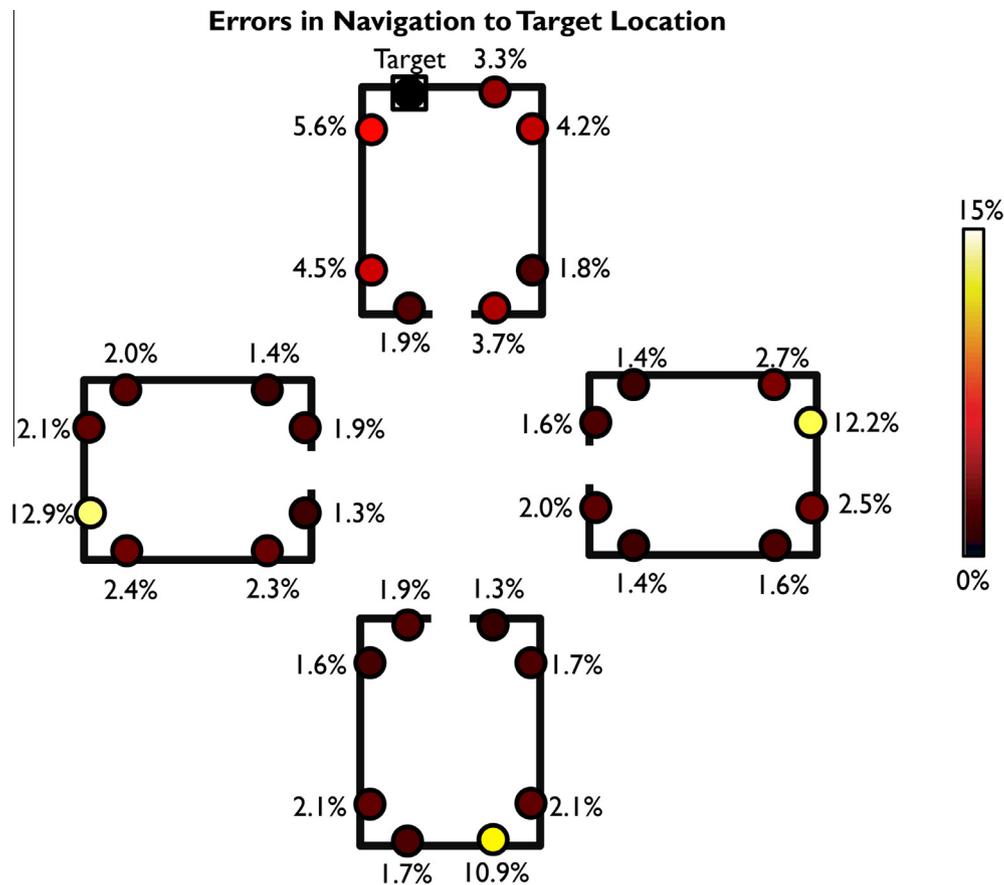


Fig. 6. Map of errors made by participants in Experiments 1–5. To create this diagram, we created an error map for each of the 32 target object locations consisting of the incorrect locations selected by participants when attempted to navigate to that target. The target object here is indicated by the black square; all incorrect locations are indicated by circles. Solid lines indicate the bounding geometry of the museums used in Experiment 1 and 2; Experiments 3–5 used different geometry but analogous locations (see Fig. 2). These error maps were aligned to one another to maintain the relative spatial relationships between the correct target location and all incorrect locations, summed together and then divided by the total number of errors. For example, the value for the location directly to the right of the target (indicated by the square) represents the percentage of errors participants made by navigating to the object along the same wall as the target. Locations were colored according to the frequency with which participants navigated to that location when attempting to find the target object. The map clearly shows evidence for schematic confusions between locations in analogous positions in different environments.

correctly identified the true location of 19.8 (SD = 8.7) of 32 goals on average, suggesting they had formed a memory representation of the environment but made a sufficient number of errors for further analysis. Two participants located every target successfully, and thus were excluded from analysis of their errors. We then classified incorrect trials as context-preserving, schema-preserving, or non-systematic errors. Schema-preserving errors were the most common (M = 45.7% of all errors, SD = 6.6%), followed by non-systematic errors (M = 27.5%, SD = 3.4%), and context-preserving errors (M = 26.8%, SD = 5.5%). The preponderance of schema-preserving errors was striking, given that these errors should have been the least common if participants were guessing locations at random. That is, of the 31 incorrect locations participants could have navigated to, only 3 were in correct local location but the wrong museum (see Fig. 1) limiting the number of ways participants could have made schema-preserving errors by chance. In contrast, 7 incorrect locations were in the same museum as the target object, and thus were potential context-preserving errors, and the remaining 21 locations were neither in the correct museum nor the correct location. In summary, schema-preserving errors were the most frequent despite being half as likely as context-preserving errors and a seventh as likely as non-systematic errors.

We then formally tested whether subjects made schema-preserving and/or context-preserving errors more often than

would be expected by chance. Indeed, participants made schema-preserving errors more often than would be expected by chance, $t(21) = 5.409$, $p < 0.001$, but context-preserving errors were at chance levels, $t(21) = 0.775$, $p = 0.447$. We then normalized each error rate by its baseline probability so that we could directly compare the rates of schema-preserving and context-preserving errors (Fig. 3). A paired-samples t -test on these values confirmed that the normalized error proportion was greater for schema-preserving errors than for context-preserving errors, $t(21) = 3.912$, $p < 0.001$. In other words, compared to what would be expected by chance, participants erred more often by navigating to the correct location in the wrong museum—even though recalling the identity of the museum would have brought their search closer to their goal. Finally, we did not analyze the frequency of non-systematic errors as this would be redundant with the previous analyses—our classification of errors is zero sum and any increase in systematic errors (either schema-preserving or context-preserving) must create a corresponding decrease in the frequency of non-systematic errors (rates of non-systematic errors for all experiments are reported in Table 1). Taken together, the results were consistent with a schematic representation of object locations that could be recalled independently of memory for the identity of the environment containing the object.

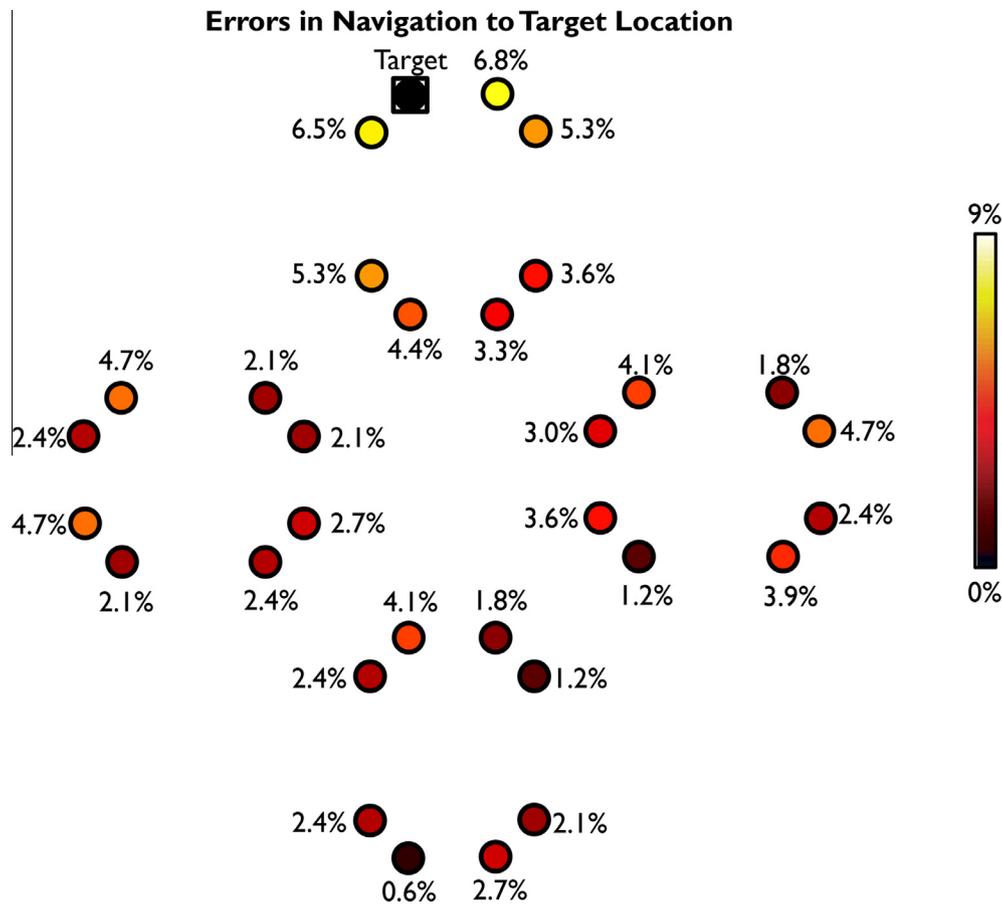


Fig. 7. Map of errors made by participants in Experiments 6 & 7. This map was constructed in the same manner as Fig. 6, but based on the data from Experiments 6 & 7 where the bounding walls of the museums were removed. Notably, participants did not make schematic errors across different groupings of objects.

Table 1

Descriptive statistics for the percentage of errors classified as each error types for all experiments (mean \pm standard deviation).

Experiment	Schema-preserving errors	Context-preserving errors	Non-systematic errors
1	45.7% \pm 6.6%	26.8% \pm 5.5%	27.5% \pm 3.4%
2	32.5% \pm 3.5%	32.0% \pm 4.3%	35.4% \pm 4.2%
3	38.8% \pm 5.1%	24.8% \pm 4.0%	36.4% \pm 4.7%
4	44.8% \pm 3.9%	25.5% \pm 3.3%	29.6% \pm 3.4%
5	35.2% \pm 5.3%	25.7% \pm 3.4%	39.1% \pm 4.8%
6	8.4% \pm 2.2%	45.2% \pm 5.8%	46.4% \pm 5.6%
7	13.6% \pm 3.4%	34.0% \pm 4.7%	52.3% \pm 5.2%

3.2. Experiment 2: Schematic representations reflect memory for spatial layout and not path of travel

Although the previous results suggest that participants used geometry-based spatial schemas to encode the locations of the target objects, an alternative possibility is that they simply memorized a fixed path from the door to the position of the target object (O'Keefe & Nadel, 1978; Packard & McGaugh, 1996). Such a strategy would give the same pattern of errors, as these paths would lead to geometrically equivalent locations in different museums. To rule out this alternative account, we ran a second experiment (Experiment 2) in which participants entered the museums through different doors during the free exploration and testing phases. The environment was the same as in Experiment 1, except that doors were added to a side wall in each museum (Fig. 2A); these side doors were blocked by barricades during the

free exploration phase. During the navigation test the barricades were moved to block the main doors, forcing participants to enter the museums through the side doors.

Participants' errors were similar to those seen in Experiment 1 (Fig. 3), even though they entered the museums through different doors during training and test. Participants correctly located 14.0 targets on average ($SD = 8.1$), and they made both schema-preserving errors ($M = 32.5\%$; $SD = 3.5\%$, $t(23) = 6.489$, $p < 0.001$) and context-preserving errors ($M = 32.0\%$; $SD = 4.3\%$, $t(23) = 2.189$, $p = 0.039$) more often than would be predicted by chance. Nevertheless, the normalized proportion of errors was significantly greater for schema-preserving than context-preserving errors, $t(23) = 4.072$, $p < 0.001$. Thus, our control experiment suggests that participants' behavior reflects the use of a schematic representation of location rather than the implementation of a simple response or rule to reach an object's location.

3.3. Experiments 3–5: Schematic representations generalize to environments with different geometries

We then considered how our results might be affected by various spatial details of our virtual environment. One aspect of the design that may have encouraged schematic coding was the fact that all of the museums were geometrically equivalent—our everyday environments are rarely so perfectly consistent. It is possible that schematic errors are found only when distinct environments can be put into perfect correspondence. We tested this in Experiment 3 by changing the geometry of two of the museums, so that two museums were circular and two were rectangular (see Fig. 2B).

If our results reflected a special case for geometrically equivalent environments, then schema-preserving errors should only be observed across museums that share the same shape. Alternatively, if the errors reflect the use of a more flexible spatial schema, then we might expect to observe schematic confusions between the rectangular and circular museum.

Participants in Experiment 3 correctly located an average of 15.1 goals ($SD = 8.3$); one participant had perfect performance and thus was excluded from analysis. On incorrect trials they made schema-preserving errors ($M = 38.8\%$; $SD = 5.1\%$) more often than would be expected by chance, $t(22) = 5.611$, $p < 0.001$, but not context-preserving errors ($M = 24.8\%$; $SD = 4.0\%$), $t(22) = 0.549$, $p = 0.589$, and the normalized proportion of errors was greater for the schema-preserving errors, $t(22) = 4.545$, $p < 0.001$. Critically, participants made schema-preserving errors both when the chosen museum had the same shape as the correct museum and also when it had a different shape. Indeed, if schema-preserving errors did not depend at all on an exact geometric match, then we would expect them to be divided evenly across the three incorrect museums, one of which has the same geometry as the correct museum (1/3 of the errors) and two of which have different geometries (2/3 of the errors). We calculated the proportion of schema-preserving errors that occurred in same and different geometry museums (excluding one additional participant who made no schema-preserving errors) and observed proportions in line with this prediction: roughly twice as many schematic errors reflected navigation to different geometry museums ($M = 66.6\%$ of all schema-preserving errors; $SD = 4.3$) than to same geometry museums ($M = 33.4\%$ of all schema-preserving errors; $SD = 4.3\%$). We explicitly tested this by doubling each participant's same-geometry error rate and comparing this to their different-geometry error rate—if the errors were evenly divided across museums then we should observe no difference. This is what we observed, $t(21) = -0.010$, $p = 0.992$. These results suggest that the same schematic representation was used across all museums, irrespective of the shape of the room.

We next tested whether the critical determinant of geometric similarity was the exact arrangement of the alcoves containing the target objects. If subjects used an intrinsic reference frame based on the relationships among the eight target objects in the environment (Mou & McNamara, 2002), then the rectangular and circular museums would have *effectively* the same geometry because the alcoves were in the same configuration in both. To rule this out, we included two variants of the experiment in which both the shape of the museums and alcove configuration differed between museums. In Experiment 4, two of the museums were squares and two were rectangles; in Experiment 5, two of the museums were the rectangles used in previous experiments and two were wider rectangles. In both experiments we again observed that participants reliably made schema-preserving errors (Exp. 4—squares vs. rectangles: correct $M = 15.2$; $SD = 8.8$ with two participants excluded for perfect performance; schema-preserving errors: $t(21) = 8.910$, $p < 0.001$; context-preserving errors: $t(21) = 0.891$, $p = 0.383$; schematic > context: $t(21) = 6.929$, $p < 0.001$; Exp. 5—rectangles with different aspect ratios: correct $M = 12.9$; $SD = 8.1$; schema-preserving errors: $t(23) = 4.744$, $p < 0.001$; context-preserving errors: $t(23) = 0.925$, $p = 0.364$; schematic > context: $t(23) = 3.914$, $p < 0.001$). Moreover, these schema-preserving errors were evenly distributed across different-geometry and same-geometry museums in both experiments (Exp. 4—squares vs. rectangles: $M = 67.9\%$ of schema-preserving errors in different geometry museums vs. $M = 32.1\%$ of schema-preserving errors in same geometry museums, $t(21) = 0.281$, $p = 0.781$; Exp. 5—rectangles with different aspect ratios: $M = 70.9\%$ of schema-preserving errors in different geometry museums vs. $M = 29.1\%$ of schema-preserving errors in same

geometry museums, $t(23) = 0.865$, $p = 0.396$). Thus, subject's schematic errors generalized across changes in bounding geometry and alcove configuration suggesting a flexible reference frame that capture the relative, rather than metric, geometry of the environment.

3.4. Experiments 6 & 7: Schematic representations require bounded environments

The fact that schema-preserving errors generalized across museums of different geometry might imply that bounding geometry is unimportant for defining spatial schemas. In Experiment 6, we tested whether the arrangement of alcoves into regular configurations (rectangles) which could be aligned to one another (Loewenstein & Gentner, 2001) was sufficient to produce schematic representations. To do this, we modified the basic virtual environment from Experiment 1 by removing all of the walls, door, and internal architectural features, but keeping the objects and their alcoves in exactly the same positions. If a regular configuration of objects was sufficient to induce schematic coding, then we might still observe schema-preserving errors even in this open-field environment.

Instead, we observed a strikingly different pattern of results (Fig. 5). Participants continued to make errors, correctly locating 20.4 out of 32 targets on average ($SD = 9.4$) with two participants exhibiting perfect performance. However, the nature of the errors changed. They no longer made schema-preserving errors at rates greater than chance ($M = 8.4\%$; $SD = 2.2\%$, $t(21) = 0.591$, $p = 0.560$). Rather, they made what appeared to be context-preserving errors ($M = 45.2\%$; $SD = 5.8\%$, $t(21) = 3.882$, $p < 0.001$), which were only observed at above-chance rates in one of the previous five experiments. In contrast to these previous experiments, here the normalized proportion of errors was greater for context-preserving errors than for schema-preserving errors, $t(21) = 2.956$, $p = 0.007$.

Thus, removing the walls changed the types of errors participants made in localizing targets. To confirm the differences between experiments, we compared systematic error rates from Experiment 6 to Experiment 1 using a 2×2 mixed-model ANOVA with a within-subject factor of error-type (schema-preserving or context-preserving) and a between-subject factor of experiment. We observed a significant main effect of experiment, $F(1,42) = 8.245$, $p = 0.049$, but no main effect of error-type, $F(1,46) = 1.795$, $p = 0.188$. Critically, we also observed a significant interaction of error-type and experiment, $F(1,46) = 17.221$, $p < 0.001$, confirming that change from schematic to contextual errors between the two experiments, which were identical except for the presence or absence of bounding walls.

In Experiment 7, we tested whether schematic errors could be restored in the open-field environment by using landmarks in the center of each cluster of objects. Our thought was that the landmarks could facilitate the anchoring of a local reference frame to each cluster, and thus substitute in part for the missing walls. Despite the presence of landmarks, we observed the same pattern as in Experiment 6 with no evidence for schema-preserving errors and a tendency to make context-preserving errors instead (correct $M = 17.9$; $SD = 9.9$; schema-preserving errors: $M = 13.6\%$; $SD = 3.4\%$, $t(23) = 1.154$, $p = 0.260$; context-preserving errors: $M = 34.0\%$; $SD = 4.7\%$, $t(23) = 2.412$, $p = 0.024$), although the direct comparison of normalized error proportions was not significant, $t(23) = 0.217$, $p = 0.30$. A 2×2 mixed model ANOVA with factors for error type and experiment compared these data to Experiment 1 and again found a significant interaction, $F(1,44) = 9.132$, $p = 0.004$, suggesting that removing the boundaries created a reliable and replicable change in participants spatial representations whether landmarks were present or not.

3.5. Summary of navigational errors in bounded and open-field environments

Finally, we created diagrams to comprehensively illustrate all errors made by participants in experiments with boundaries (Experiments 1–5, Fig. 6) and without (Experiments 6–7, Fig. 7). To create these diagrams, we created an error map for each of the 32 target object locations consisting of the incorrect locations selected by participants when attempted to navigate to that specific target. These 32 error maps were then aligned to one another to maintain the relative spatial relationships between the correct target location and all incorrect locations and summed together. The resulting map for nested environments clearly shows evidence for schematic confusions between locations in analogous positions. In contrast the map for open-field environments shows that errors are primarily induced by spatial proximity to the target in this case.

We also investigated whether there was a relationship between the tendency to make systematic errors and overall performance. We observed a significant correlation between the proportion of schema-preserving errors made on incorrect trials and overall performance in Experiments 1–5, $r(113) = 0.286$, $p = 0.002$, and we observed a similar relationship for context-preserving errors, $r(113) = 0.212$, $p = 0.023$. Thus, participants who made systematic errors rather than random errors were also more likely to correctly locate targets. This suggests that these errors reflect an intermediate stage of learning in which subjects have learned the locations of objects within each local environment and have learned the identities of the environments, but have not fully integrated this information.

4. Discussion

Our results demonstrate that humans use schematic representations of local space to recall object locations during navigation. In an environment consisting of four museums, participants frequently navigated to the correct location as defined by the local geometry of each museum, but to the wrong museum. Thus, they remembered where the object was within the local spatial reference frame, but they did not remember the identity of the context to which the frame should be applied. This pattern of results, which we replicated in five experiments, demonstrates that spatial memory retrieval does not always proceed in a coarse-to-fine manner; rather, local and global spatial information can be accessed in parallel and can occasionally be mismatched. Notably, schematic errors were *not* observed when the walls of the museums were removed in Experiments 6 and 7; this shows that navigational boundaries are crucial for delineating the extent (and possibly the alignment) of the local spatial frame.

Together, our results provide evidence for context-independent spatial reference frames that are anchored by local geometry and can be reapplied to multiple environmental contexts. We use the term “schematic” to refer to these reference frames because they appear to represent the essential spatial structure of each environment, without including details that allow spatially similar environments to be distinguished from one another. Indeed, like a schematic drawing, they allow the system to draw analogies between different local environments that are spatially similar to each other. Although the terms “schematic” and “schema” have been widely used and thus have acquired a diffuse definition, at their core these terms refer to a representation that captures the essential structure of an event or situation and specifies how elements of that situation relate to one another—for example, a “script” for the typical sequences of actions that make up an event or narrative (Bartlett, 1932; Rumelhart, 1980; Schank & Abelson,

2013). Crucially, schemas can be reused as a framework for encoding new situations that closely fit within their structure (Bartlett, 1932; Brewer & Treyns, 1981). Our participants’ confusions across museums suggest exactly such an application of the same representational framework to different situations that are structurally similar to each other. Moreover, the idea that a local spatial reference frame may serve as a spatial schema is consistent with recent evidence demonstrating that pre-existing spatial memories can allow rapid encoding (McKenzie et al., 2014; Tse et al., 2007; van Buuren et al., 2014) and consolidation (Tse et al., 2007, 2011) of novel object locations encountered within a familiar spatial environment.

4.1. Separate representations for spatial reference and environmental identity

The idea that there are separate representations for defining coordinates in the local environment and identifying which local environment one is in has been anticipated by previous results in the literature. For example, recent work in our lab demonstrated a behavioral dissociation between local coordinate and global identity representations when rodent navigators were required to reorient themselves to their surroundings before locating a goal (Julian, Keinath, Muzzio, & Epstein, 2015). Previous studies had established that after disorientation, rodents (Cheng, 1986), humans (Hermer & Spelke, 1994; Learmonth, Newcombe, & Huttenlocher, 2001; Lee & Spelke, 2011), and many other species (Lee, Ferrari, Vallortigara, & Sovrano, 2015; Vallortigara, 2009) exhibit “geometric errors” in their search patterns. For example, in a rectangular experimental chamber with a hidden goal in one of the corners, they will search equally often at the correct corner and the geometrically equivalent corner that is diagonally opposite. Notably, this search pattern is often unaffected by non-geometric cues such as colors or patterns on a wall that could potentially disambiguate these two geometrically equivalent corners. We demonstrated that mice will, however, use these non-geometric cues to determine which chamber they are in—even while simultaneously ignoring them as a guide to locations within the chamber (Julian et al., 2015). That is, the mice used visual features to establish their coarse location within the wider world (e.g. Box A) while failing to use those features to recover the precise location of their goal within that context. This dissociation indicates that—in mice at least—identification of the local environment and spatial localization within that environment are independent processes that most likely rely upon different underlying representations.

Previous neuroimaging and neuropsychological results also point to a similar distinction in humans between neural mechanisms for place/context recognition, which are believed to be primarily supported by the parahippocampal place area (PPA; Aguirre, Zarahn, & D’Esposito, 1998; Epstein & Kanwisher, 1998; Epstein & Vass, 2014; Marchette, Vass, Ryan, & Epstein, 2015), and neural mechanisms for spatial orientation, which are believed to be primarily supported by the retrosplenial complex (RSC; Baumann & Mattingley, 2010; Epstein, 2008; Marchette et al., 2014; Vass & Epstein, 2013). Patients with damage to the lingual/parahippocampal region that encompasses PPA cannot visually identify familiar surroundings, such as buildings or oft-traveled intersections, but can draw maps of well-known environments (Epstein, DeYoe, Press, Rosen, & Kanwisher, 2001; Mendez & Chierri, 2003). In contrast, patients with damage to medial parietal cortex, including RSC, can recognize their current environment but cannot use this information to orient themselves to it or guide navigation (Aguirre & D’Esposito, 1999; Hashimoto, Tanaka, & Nakano, 2010; Takahashi & Kawamura, 2002). This neural distinction between place/context recognition and spatial orientation

anticipates the current results, where a similar distinction is observed in behavior.

It is worth noting, however, that the current results do not imply that spatial reference frames are always schematic, or local. Although the reliance on local reference frames in the current study is consistent with previous observations that people are inaccurate when pointing to objects outside of their current environment (Montello, 1991; Wang & Brockmole, 2003a, 2003b) and is predicted by a recent model of spatial memory (Meilinger, 2008), other data suggest that we may sometimes form global representations that exist in parallel with local representations (Greenauer & Waller, 2010; Meilinger, Riecke, & Bühlhoff, 2014). The ability to integrate local divisions into a global representation may require more extensive experience than was provided to our subjects, or particularly strong navigation abilities (Ishikawa & Montello, 2006; Schinazi & Epstein, 2010; Weisberg, Schinazi, Newcombe, Shipley, & Epstein, 2014). An intriguing possibility is that spatial representations are initially generated in schematic form, but over time come to include episodic or contextual details that allow local environments to be distinguished from one another and integrated into a larger structure (Carpenter, Manson, Jeffery, Burgess, & Barry, 2015; Lever, Wills, Cacucci, Burgess, & O'Keefe, 2002; Uttal, Friedman, Hand, & Warren, 2010). Consolidated topographical knowledge, then, might include representations that bind together schematic location information with environmental context information.

4.2. The importance of boundaries

A striking aspect of our data was that schematic errors were only observed at above-chance rates when the museums were defined by bounding walls. When the walls were removed, the number of schematic errors returned to chance, despite the fact that the arrangement of the objects, and the alcoves that they rested upon, remained unchanged. This result suggests that the bounding walls were crucial for subdividing the larger environment into local spaces and for anchoring the local reference frame used to define locations within these spaces.

These results are consistent with previous work indicating that bounding geometry is an important cue for encoding object locations (Doeller, King, & Burgess, 2008; Hartley, Trinkler, & Burgess, 2004). Indeed, “border” and “boundary” cells, which track an animal's distance and direction to nearby boundaries, have been identified in the rodent hippocampal formation (Lever, Burton, Jeewajee, O'Keefe, & Burgess, 2009; O'Keefe & Burgess, 1996; Solstad, Boccara, Kropff, Moser, & Moser, 2008). Moreover, in the fMRI experiment upon which the current study is based, we found that RSC coded headings and locations relative to local boundaries (Marchette et al., 2014). Behavioral studies have found that merely passing through a doorway impairs memory for the identity and position of objects encountered in the previous room, suggesting that boundaries play an important role in segmenting memory representations (Horner, Bisby, Wang, Bogus, & Burgess, 2016; Radvansky & Copeland, 2006). By dividing the larger environments into a number of smaller local divisions, the bounding walls in the current study would have facilitated the encoding of hierarchical spatial memories in which each object's location was simultaneously represented at two levels—first, the identity of the local context (or “category” Huttenlocher et al., 1991) containing the object; and second, its metric position within that local context (Huttenlocher et al., 1994; Kosslyn et al., 1974; McNamara et al., 1989; Meilinger, 2008; Stevens & Coupe, 1978). Once the division into different local environments was established, then geometric similarities between the museums would have made it possible for them to be aligned to each other (DeLoache, 1987; Loewenstein & Gentner, 2001) and the same spatial schema to be

applied to each (Bartlett, 1932; Piaget, 1929; Tse et al., 2007). In contrast, when the boundaries were removed in Experiments 6 and 7, the environment was no longer subdivided, and the “local” reference frame might have expanded to include all of the objects within the park (c.f., Li & Gleitman, 2002). Indeed, participants in these experiments no longer made schematic errors but rather navigated to the correct cluster of objects (where a museum would be), while occasionally misremembering the true location within it.

Although boundaries may have encouraged schematic coding, the results of Experiments 3–5 indicate that an exact geometric match was not required for participants to apply a schematic reference frame across different environments. Our participants confused analogous locations even across museums with different shapes and aspect ratios. What kind of representation might have allowed subjects to draw these analogies while tolerating discrepancies in the geometry of the walls and alcoves?

We can suggest three possibilities. First, subject might have coded the route from the door to the object in addition to coding the location of the object in the museum (e.g., Alexander & Nitz, 2015). This might have led to confusions across locations in different-shaped museums because of route similarity. Although the results of Experiment 2, in which different entrances were used at study and test, would seem to argue against route coding, in fact these results only indicate that confusions cannot be *exclusively* attributed to route coding; they do not exclude the possibility that locations were coded using routes and geometry simultaneously. Second, participants might have coded locations relative to geometry, but in an imprecise manner that led to confusions between locations that had similar but not precisely identical bearings and distances to the walls. Supporting this idea, models of location coding that depend upon bearings to geometry can generalize across changes in environmental size and aspect ratio (Barry et al., 2006; Hartley et al., 2004; O'Keefe & Burgess, 1996). Third, participants may have oriented themselves within the museum based upon the position of the main entryway and their orientation upon passing through it. This idea is supported by observations that reference frames can be oriented by many factors, including the environmental/intrinsic geometry of the environment (Mou & McNamara, 2002; Shelton & McNamara, 1997), and also one's egocentric experience within the environment (Greenauer & Waller, 2008; Shelton & McNamara, 1997, 2001). In this account, boundaries would have been used in the current experiment to segment the environment into areas amenable to local coding, while the orientation of these local spatial reference frames would have been based on the principal axis of the environment defined by the door and the axis of the room.

4.3. Why two mechanisms?

Why would our cognitive architecture include separate representations for local space and contextual identity, which can potentially become incorrectly matched to each other? We conjecture that this division may relate to the challenges that the navigational system faces when anchoring spatial codes to the world. For mental spatial coordinates to meaningfully guide behavior, they must be referenced to fixed elements of the external environment. For example, GPS coordinates reveal an object's location on the earth's surface, typically using the equator and prime meridian as reference points. In the absence of satellite telemetry, however, it is difficult to solve this referencing problem, as there are no perceptible features that allow one to uniquely locate one's position across all terrestrial environments. Our navigational systems might solve this referencing problem by using features within the local (i.e. currently visible) environment—such as walls or other landmarks—to anchor spatial coordinates onto the external world

(Marchette et al., 2014). However, this still leaves the challenge of relating these local coordinates to the spatial structure of the wider world.

A critical part of solving the second part of this problem would be the ability to distinguish one local environment from another, even if they have similar intrinsic structure. Although one might theoretically use the same mechanism to distinguish between environments that one uses to define coordinates within each environment, the features that are most useful for the former task might not be the same features that are most useful for the latter. In particular, whereas the shape of the environment as defined by geometric boundaries might be critical for identifying locations (and headings) within an environment, there might not be enough possible variations in geometry to distinguish between all of the navigational contexts that we encounter in our lives. Thus, there would be a need for two kinds of representation: a mostly geometric representation that is used to define locations and headings in local space, and a more varied representation that uses a wide range of potential landmark cues (such as geometric layout, non-geometric visual features, and object identities) to distinguish between different places or navigational contexts.

5. Summary

Our results suggest that human spatial memory may contain two dissociable representations of where an object can be found in nested environments—a schematic map of where an object lies with respect to local geometry, and a representation of the identity and location of the environmental context. Boundaries and the passages through them appear to play a special role in segmenting a continuous navigable space into local environmental contexts to which schematic representations of location may be applied. When two such local regions have similar configurations, a single schematic reference frame may be used to encode locations in both environments, creating spatial analogies between objects in different places in the world. By taking advantage of these analogies, we can reuse our cognitive maps, thus allowing memory for past explorations to serve as a scaffold for encoding the new places we encounter in our travels.

Acknowledgements

This research was supported the National Institutes of Health (EY022350) and the National Science Foundation (SBE-0541957, SBE-1041707). We would like to thank Rebecca Somach for assistance with data collection.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.cognition.2016.10.005>.

References

- Aguirre, G. K., Zarahn, E., & D'Esposito, M. (1998). An area within human ventral cortex sensitive to "building" stimuli: evidence and implications. *Neuron*, *21*(2), 373–383.
- Alexander, A. S., & Nitz, D. A. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. *Nature Neuroscience*, *18*(8), 1143–1151.
- Barry, C., Lever, C., Hayman, R., Hartley, T., Burton, S., O'Keefe, J., ... Burgess, N. (2006). The boundary vector cell model of place cell firing and spatial memory. *Reviews in the Neurosciences*, *17*(1–2), 71–98.
- Bartlett, F. C. (1932). *Remembering: An experimental and social study*. Cambridge: Cambridge University.
- Baumann, O., & Mattingley, J. B. (2010). Medial parietal cortex encodes perceived heading direction in humans. *The Journal of Neuroscience*, *30*(39), 12897–12901.
- Brewer, W. F., & Treyens, J. C. (1981). Role of schemata in memory for places. *Cognitive Psychology*, *13*(2), 207–230.
- Carpenter, F., Manson, D., Jeffery, K., Burgess, N., & Barry, C. (2015). Grid cells form a global representation of connected environments. *Current Biology*, *25*(9), 1176–1182.
- Cheng, K. (1986). A purely geometric module in the rat's spatial representation. *Cognition*, *23*(2), 149–178.
- Colgin, L. L., Moser, E. I., & Moser, M.-B. (2008). Understanding memory through hippocampal remapping. *Trends in Neurosciences*, *31*(9), 469–477.
- DeLoache, J. S. (1987). Rapid change in the symbolic functioning of very young children. *Science*, *238*(4833), 1556–1557.
- Doeller, C. F., King, J. A., & Burgess, N. (2008). Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *Proceedings of the National Academy of Sciences*, *105*(15), 5915–5920.
- Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature*, *425*(6954), 184–188.
- Epstein, R. A. (2008). Parahippocampal and retrosplenial contributions to human spatial navigation. *Trends in Cognitive Sciences*, *12*(10), 388–396.
- Epstein, R. A., DeYoe, E. A., Press, D. Z., Rosen, A. C., & Kanwisher, N. (2001). Neuropsychological evidence for a topographical learning mechanism in parahippocampal cortex. *Cognitive Neuropsychology*, *18*(6), 481–508.
- Epstein, R. A., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*(6676), 598–601.
- Epstein, R. A., & Vass, L. K. (2014). Neural systems for landmark-based wayfinding in humans. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *369*(1635), 20120533.
- Greenauer, N., & Waller, D. (2008). Intrinsic array structure is neither necessary nor sufficient for nonegocentric coding of spatial layouts. *Psychonomic Bulletin & Review*, *15*(5), 1015–1021.
- Greenauer, N., & Waller, D. (2010). Micro-and macrorference frames: Specifying the relations between spatial categories in memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*(4), 938.
- Hartley, T., Trinkl, I., & Burgess, N. (2004). Geometric determinants of human spatial memory. *Cognition*, *94*(1), 39–75.
- Hashimoto, R., Tanaka, Y., & Nakano, I. (2010). Heading disorientation: A new test and a possible underlying mechanism. *European Neurology*, *63*(2), 87–93.
- Hermer, L., & Spelke, E. S. (1994). A geometric process for spatial reorientation in young children. *Nature*.
- Hirtle, S. C., & Jonides, J. (1985). Evidence of hierarchies in cognitive maps. *Memory & Cognition*, *13*(3), 208–217.
- Horner, A. J., Bisby, J. A., Wang, A., Bogus, K., & Burgess, N. (2016). The role of spatial boundaries in shaping long-term event representations. *Cognition*, *154*, 151–164.
- Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, *98*(3), 352.
- Huttenlocher, J., Newcombe, N., & Sandberg, E. H. (1994). The coding of spatial location in young children. *Cognitive Psychology*, *27*(2), 115–147.
- 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.
- Jeffery, K. J., & Anderson, M. I. (2003). Dissociation of the geometric and contextual influences on place cells. *Hippocampus*, *13*(7), 868–872.
- Julian, J. B., Keinath, A. T., Muzzio, I. A., & Epstein, R. A. (2015). Place recognition and heading retrieval are mediated by dissociable cognitive systems in mice. *Proceedings of the National Academy of Sciences*, *112*(20), 6503–6508.
- Kosslyn, S. M., Pick, H. L., & Fariello, G. R. (1974). Cognitive maps in children and men. *Child Development*, *45*, 707–716.
- Learmonth, A. E., Newcombe, N. S., & Huttenlocher, J. (2001). Toddlers' use of metric information and landmarks to reorient. *Journal of experimental child psychology*, *80*(3), 225–244.
- Lee, S. A., Ferrari, A., Vallortigara, G., & Sovrano, V. A. (2015). Boundary primacy in spatial mapping: Evidence from zebrafish (*Danio rerio*). *Behavioural Processes*, *119*, 116–122.
- Lee, S. A., & Spelke, E. S. (2011). Young children reorient by computing layout geometry, not by matching images of the environment. *Psychonomic Bulletin & Review*, *18*(1), 192–198.
- Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., & Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *The Journal of Neuroscience*, *29*(31), 9771–9777.
- Lever, C., Wills, T., Cacucci, F., Burgess, N., & O'Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature*, *416*(6876), 90–94.
- Li, P., & Gleitman, L. (2002). Turning the tables: Language and spatial reasoning. *Cognition*, *83*(3), 265–294.
- Loewenstein, J., & Gentner, D. (2001). Spatial mapping in preschoolers: Close comparisons facilitate far mappings. *Journal of Cognition and Development*, *2*(2), 189–219.
- Marchette, S. A., Vass, L. K., Ryan, J., & Epstein, R. A. (2014). Anchoring the neural compass: Coding of local spatial reference frames in human medial parietal lobe. *Nature Neuroscience*, *17*(11), 1598–1606.
- Marchette, S. A., Vass, L. K., Ryan, J., & Epstein, R. A. (2015). Outside looking in: Landmark generalization in the human navigational system. *The Journal of Neuroscience*, *35*(44), 14896–14908.
- McKenzie, S., Frank, A. J., Kinsky, N. R., Porter, B., Rivière, P. D., & Eichenbaum, H. (2014). Hippocampal representation of related and opposing memories develop

- within distinct, hierarchically organized neural schemas. *Neuron*, 83(1), 202–215.
- McNamara, T. P., Hardy, J. K., & Hirtle, S. C. (1989). Subjective hierarchies in spatial memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(2), 211.
- Meilinger, T. (2008). *The network of reference frames theory: A synthesis of graphs and cognitive maps spatial cognition VI. Learning, reasoning, and talking about space*. Springer.
- Meilinger, T., Riecke, B. E., & Bühlhoff, H. H. (2014). Local and global reference frames for environmental spaces. *The Quarterly Journal of Experimental Psychology*, 67(3), 542–569.
- Mendez, M. F., & Chierri, M. M. (2003). Agnosia for scenes in topographagnosia. *Neuropsychologia*, 41(10), 1387–1395.
- Miller, J. F., Neufang, M., Solway, A., Brandt, A., Trippel, M., Mader, I., ... Jacobs, J. (2013). Neural activity in human hippocampal formation reveals the spatial context of retrieved memories. *Science*, 342(6162), 1111–1114.
- Montello, D. R. (1991). Spatial orientation and the angularity of urban routes a field study. *Environment and Behavior*, 23(1), 47–69.
- Montello, D. R. (1993). *Scale and multiple psychologies of space spatial information theory a theoretical basis for GIS*. Springer.
- Mou, W., & McNamara, T. P. (2002). Intrinsic frames of reference in spatial memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(1), 162.
- Muller, R. U., & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *The Journal of Neuroscience*, 7(7), 1951–1968.
- Newcombe, N., Huttenlocher, J., Sandberg, E., Lie, E., & Johnson, S. (1999). What do misestimations and asymmetries in spatial judgement indicate about spatial representation? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(4), 986.
- O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381(6581), 425–428.
- 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.
- O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map* (Vol. 3) Oxford: Clarendon Press.
- Packard, M. G., & McGaugh, J. L. (1996). Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiology of Learning and Memory*, 65(1), 65–72.
- Aguirre, G. K., & D'Esposito, M. (1999). Topographical disorientation: A synthesis and taxonomy. *Brain*, 122(9), 1613–1628.
- Piaget, J. (1929). *The child's conception of the world*. London: Kegan Paul.
- Radvansky, G. A., & Copeland, D. E. (2006). Walking through doorways causes forgetting: Situation models and experienced space. *Memory & cognition*, 34(5), 1150–1156.
- Rumelhart, D. E. (1980). Schemata: The building blocks of cognition. In R. Spiro, B. Bruce, & W. Brewer (Eds.), *Theoretical issues in reading comprehension*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Schank, R. C., & Abelson, R. P. (2013). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Psychology Press.
- Schinazi, V. R., & Epstein, R. A. (2010). Neural correlates of real-world route learning. *Neuroimage*, 53(2), 725–735.
- Shelton, A. L., & McNamara, T. P. (1997). Multiple views of spatial memory. *Psychonomic Bulletin & Review*, 4(1), 102–106.
- Shelton, A. L., & McNamara, T. P. (2001). Systems of spatial reference in human memory. *Cognitive Psychology*, 43(4), 274–310.
- Solstad, T., Boccara, C. N., Kropff, E., Moser, M.-B., & Moser, E. I. (2008). Representation of geometric borders in the entorhinal cortex. *Science*, 322(5909), 1865–1868.
- Spiers, H. J., Hayman, R. M., Jovalekic, A., Marozzi, E., & Jeffery, K. J. (2015). Place field repetition and purely local remapping in a multicompartiment environment. *Cerebral Cortex*, 25(1), 10–25.
- Stevens, A., & Coupe, P. (1978). Distortions in judged spatial relations. *Cognitive Psychology*, 10(4), 422–437.
- Takahashi, N., & Kawamura, M. (2002). Pure topographical disorientation—The anatomical basis of landmark agnosia. *Cortex*, 38(5), 717–725.
- Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., ... Morris, R. G. (2007). Schemas and memory consolidation. *Science*, 316(5821), 76–82.
- Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., ... Morris, R. G. (2011). Schema-dependent gene activation and memory encoding in neocortex. *Science*, 333(6044), 891–895.
- Uttal, D. H., Friedman, A., Hand, L. L., & Warren, C. (2010). Learning fine-grained and category information in navigable real-world space. *Memory & Cognition*, 38(8), 1026–1040.
- Vallortigara, G. (2009). Animals as natural geometers. In L. Tommasi, L. Nadel, & M. Peterson (Eds.), *Cognitive biology: Evolutionary and developmental perspectives on mind, brain and behaviour* (pp. 83–104).
- van Buuren, M., Kroes, M. C., Wagner, I. C., Genzel, L., Morris, R. G., & Fernández, G. (2014). Initial investigation of the effects of an experimentally learned schema on spatial associative memory in humans. *The Journal of Neuroscience*, 34(50), 16662–16670.
- Vass, L. K., & Epstein, R. A. (2013). Abstract representations of location and facing direction in the human brain. *The Journal of Neuroscience*, 33(14), 6133–6142.
- Wang, R. F., & Brockmole, J. R. (2003a). Human navigation in nested environments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(3), 398.
- Wang, R. F., & Brockmole, J. R. (2003b). Simultaneous spatial updating in nested environments. *Psychonomic Bulletin & Review*, 10(4), 981–986.
- 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.
- Wills, T. J., Lever, C., Cacucci, F., Burgess, N., & O'Keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. *Science*, 308(5723), 873–876.