UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

All Eyes on the Hippocampus: The Primate Hippocampus as a Visually-Guided Cognitive Graph

Permalink

https://escholarship.org/uc/item/8444g7k8

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 47(0)

Authors

Huffman, Derek J.
Dang, Linh Khanh
Doherty, Olivia Roan
et al.

Publication Date

2025

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

All Eyes on the Hippocampus: The Primate Hippocampus as a Visually-Guided Cognitive Graph

Derek J. Huffman*, **Linh K. Dang, Olivia R. Doherty, Amrit Shakya, and Ainsley K. Bonin**
Department of Psychology, Colby College
*Corresponding author: derek.huffman@colby.edu

Abstract

Spatial memory is a core cognitive function of many mobile animals. The study of spatial cognition is highly interdisciplinary and different approaches have led to discordant hypotheses about spatial memory. The neuroscientific discovery of place cells and grid cells in the rodent hippocampus and entorhinal cortex led to the highly influential "cognitive map" hypothesis. Conversely, behavioral evidence led to the "cognitive graph" hypothesis that suggests memory is distortion-prone. We argue that between-species differences in sensory and perceptual systems cause vision to play a predominant role in primates. We build on previous modeling work by developing a visual version of the SR model to show that it can provide a unified framework to account for seemingly disparate findings from the brain and behavior, thus providing evidence for our hypothesis that primate (e.g., human) spatial cognition is driven by the hippocampal system, which instantiates a visually-guided cognitive graph.

Keywords: cognitive map; cognitive graph; place cells; grid cells; spatial view cells; hippocampus

Introduction

Given the multidisciplinary nature of the study of spatial memory, researchers have approached the study of spatial cognition with a wide variety of tools, ranging from fields as broad as geography to neuroscience. Here, we focus on approaches at the intersection of the study of neuroscience, behavior, and computational modeling, which have all offered important clues into the nature of spatial memory. Neuroscientific studies have revealed several kinds of spatially tuned cells. For example, O'Keefe and Dostrovsky (1971) discovered hippocampal "place cells", which increase their firing rate when an animal is in a specific part of the environment. Additionally, Hafting, Fyhn, Molden, Moser, and Moser (2005) discovered "grid cells" cells in the medial entorhinal cortex, which fire along the vertices of equilateral triangles, forming hexagon-shaped patterns. Influential theories used these findings to posit that spatial representations are supported by metric, Euclidean knowledge, termed the "cognitive map" hypothesis (e.g., Bellmund, Gärdenfors, Moser, & Doeller, 2018; Byrne, Becker, & Burgess, 2007; Gallistel, 1990; McNaughton, Battaglia, Jensen, Moser, & Moser, 2006; Moser & Moser, 2008; O'Keefe & Nadel, 1978).

Recent evidence has suggested that primate navigation may be heavily supported by vision. Primates have high-acuity vision, thus enabling the visual exploration of space, rather than having to physically traverse the space to build up a memory of the relative location of landmarks. In contrast, most rodents, including rats, have low acuity vision (e.g., Wallace et al., 2013); thus, some theories have suggested that olfactory and tactile cues might be more salient relevant for rodent navigation (vs. primates; e.g., Ekstrom, 2015; Rolls, 2023).

Consistent with between-species differences in perceptual systems, recent work has revealed that the canonical spatial cells observed in rodents might map onto visually-guided spatial cells in primates. For example, research has found evidence for spatial view cells, which fire when a monkey is looking at a salient object or landmark in the environment (e.g., Rolls & O'Mara, 1995; Rolls, Robertson, & Georges-François, 1997; Rolls, 2023). Thus, in contrast to the canonical place cell coding in rodents, spatial view cells respond to the location of landmarks "out there" in the visual world (Rolls, 2023). Similarly, Killian, Jutras, and Buffalo (2012) discovered evidence for visual grid cells, which fire based on the location of the monkey's fixation on the computer screen (for similar findings in humans with fMRI see: Julian, Keinath, Frazzetta, & Epstein, 2018; Nau, Navarro Schröder, Bellmund, & Doeller, 2018). Killian, Potter, and Buffalo (2015) revealed evidence for "saccade direction" cells, thus providing a visual counterpart to rodent head-direction cells (e.g., Taube, Muller, & Ranck, 1990). Altogether, recent theories have argued that vision may be fundamental to how primates represent space and the findings from neuroscience indicate that the hippocampal system in primates maps space in a similar, but complementary manner to what is observed in rodents (Ekstrom, 2015; Nau, Julian, & Doeller, 2018; Rolls, 2023). We explore the exciting possibility that reconceptualizing hippocampal networks in the primate as being heavily driven by the visual system can allow us to account for a host of neuroscientific and behavioral findings.

Behavioral approaches revealed evidence that spatial memory is better explained by non-Euclidean representations, termed the "cognitive graph" hypothesis. For example, several behavioral studies have suggested that judgments of distance can be asymmetric: i.e., AB $\neq$ BA, thus violating Euclidean postulates (e.g., Burroughs & Sadalla, 1979; Cadwallader, 1979; McNamara, 1991; Mcnamara & Diwadkar, 1997). Likewise, boundaries can exhibit strong influences on spatial memory (e.g., Cohen, Baldwin, & Sherman, 1978; Kosslyn, Pick, & Fariello, 1974; McNamara, 1986; Newcombe & Liben, 1982; also see: Hirtle & Jonides, 1985; Mc-

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society* ©2025 the author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

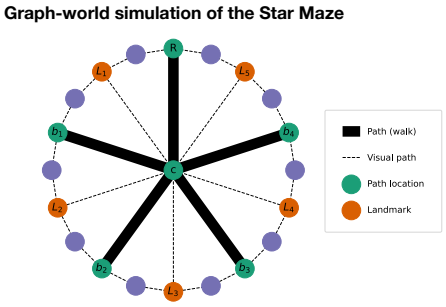


Figure 1: We modeled the star-maze task using the SR model. All of the green nodes were along the paths that the monkey could traverse whereas the other color nodes are those related to rotations and visual cues. The starting locations are indicated as b_1, b_2, b_3, and b_4; the visual landmarks are shown in orange (i.e., the L_i's), and the reward location is labeled as R.

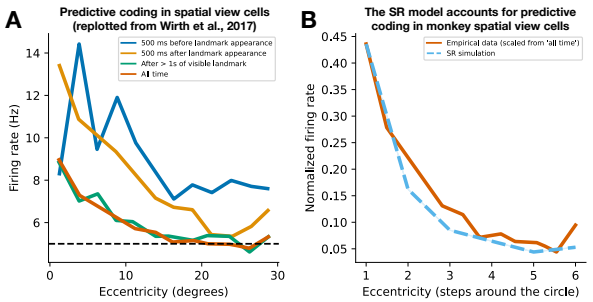


Figure 2: The SR model can account for findings from spatial view cells in the monkey hippocampus. (A) Empirical data from the monkey hippocampus found that spatial view cells exhibited predictive coding. (B) The SR spatial view cell receptive field accounts for the predictive coding in monkey spatial view cells. We scaled the empirical data to fit the the min and max values from the SR data.

Namara, Hardy, & Hirtle, 1989). Moreover, recent work with non-Euclidean mazes (e.g., that incorporate "wormholes" and other forms of impossible mazes) have shown that participants erroneously incorporate these features into their memory while also often failing to notice anything was amiss (e.g., Warren, 2019; Warren, Rothman, Schnapp, & Ericson, 2017). Altogether, these findings led to the theory that human spatial memory is error-prone and better explained by invoking the concept of a "cognitive graph" or other similar metaphors (e.g., Chrastil & Warren, 2014; McNamara, 1986, 1991; Stevens & Coupe, 1978; Tversky, 1992; Warren, 2019; Warren et al., 2017).

In summary, neuroscientific and behavioral research have developed discordant predictions about the nature of spatial representations. A recent theory attempted to reconcile these differences by positing that humans can employ both the "cognitive map" and "cognitive graph", which are supported by complementary brain networks (Peer, Brunec, Newcombe, & Epstein, 2021). Here, we explore a different, novel hypothesis that both the hippocampal network and behavior can be accounted for by the cognitive graph hypothesis. We provide support for these ideas using the Successor Representation (SR) model (Stachenfeld, Botvinick, & Gershman, 2017).

Methods

For the majority of our approach here, we followed the framework of Stachenfeld et al. (2017). Specifically, we calculated the SR matrix using a combination of the analytical method and the simulation method. For the analytical method, we calculated the SR with the following equation:

$$M = (I - \gamma T)^{-1} \tag{1}$$

where M is the SR matrix, I is the identity matrix, γ is the temporal discount factor (we list the γ for each simulation below; we note that our choice in γ scaled with the relative size of the environments below to allow for greater "blending" of states in larger environments), T is the transition matrix (each entry of the matrix indicates the probability of transitioning from one state to another state; the rows sum to 1), and $()^{-1}$ indicates matrix inversion. We also used the temporal difference rule in some simulations:

$$M[s,:] = M[s,:] + \alpha(I[s,:] + \gamma M[s_{new},:] - M[s,:]) \tag{2}$$

where α is the learning rate, I is the identity matrix, s is the state at time t and s_{new} is the state at time $t+1$.

To simulate navigation within the environment, we used a value function and a one-step look-ahead model (Russek, Momennejad, Botvinick, Gershman, & Daw, 2017). Specifically, to calculate value, we used the following equation (Stachenfeld et al., 2017):

$$V(s) = \sum_{s'} M[s,s']R[s'] \tag{3}$$

where $V(s)$ is a scalar indicating the value of state s and R is a vector of reward, which we set to a value of 1 for whatever the goal and 0 for everywhere else. The one-step look-ahead rule simply picks the adjacent state with the highest value.

To simulate visual grid cells in the primate entorhinal cortex (Killian et al., 2012), we calculated the eigendecomposition of the SR matrix using a combination of custom-written code along with code from Stachenfeld et al. (2017) (`https://github.com/kstach01/predictive_hc.git`). Briefly, we first calculated the SR matrix using the analytical method. Then, we calculated the eigenvector decomposition (using numpy.linalg.eig). We then thresholded the resultant eigenvectors at 0 by conducting element-wise multiplication with a mask of 1s wherever the eigenvalues are greater than 0 and 0s otherwise (i.e., to simulate only positive firing rates of neurons; Stachenfeld et al., 2017). We modeled the visual environment using a graph via a triangular lattice grid (using the networkx package in Python).

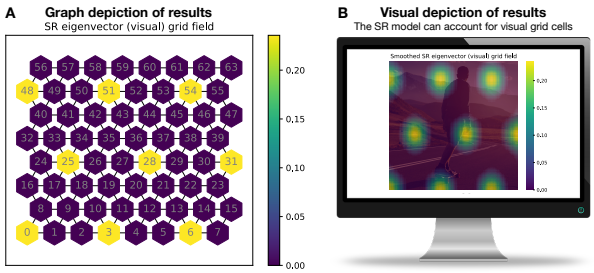


Figure 3: Eigenvector decomposition of the SR matrix can account for visual grid cells that have been observed in primates. A) A depiction of an SR eigenvector visual grid field. Here, the task structure, and hence the underlying representation, is a graph. Note the perfect hexagon and equilateral spacing of the nodes. B) We interpolated the data from the graph into a more continuous space and overlaid it on an image for a visual depiction of visual grid cells.

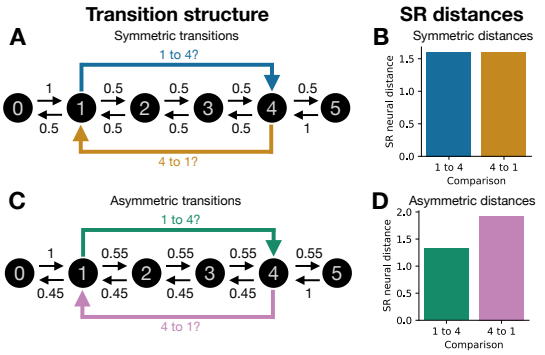


Figure 4: The SR model can account for asymmetric distance judgments. In the symmetric condition, the SR model predicts symmetric distance judgments (panels A and B). Conversely, in the asymmetric condition, the SR model predicts asymmetric distance judgments (panel C and D), thus accounting for extant findings in the behavioral literature.

Results and Discussion

Given the apparent disconnect between theories that have been advanced on the basis of neuroscientific findings vs. behavioral findings, we aimed to model findings from both neuroscience and behavior to try to better understand the possible (computational) link between these levels of study.

The SR Model Accounts for Findings from Spatial View Cells and Spatial Grid Cells in Primates

The SR Model Accounts for Spatial View Cells Spatial view cells have been observed in the hippocampus across several studies in nonhuman primates (see Introduction). Wirth, Baraduc, Planté, Pinède, and Duhamel (2017) recently studied spatial view cells in non-human primates as they performed a memory task in a virtual environment. Specifically, they trained monkeys to navigate, via a joystick, to a goal location from variable starting locations within a star maze that contained visual cues around the outer part of the environment. They found that monkeys could rapidly learn to navigate to the goal location, often requiring only approximately 10 trials to achieve $> 90\%$ accuracy. We implemented an SR simulation of their behavioral task by creating a graph of their task states (Figure 1). After training ($\gamma = 0.7$; we calculated the SR via the analytical method), we found that the SR took the shortest path to the goal state from each of the 4 starting locations. Therefore, the SR model can account for the behavioral findings in monkeys.

Wirth et al. (2017) provided evidence for what they called "anticipatory" firing of spatial view cells as a function of the eccentricity of the landmark from foveation (measured in degrees; see replotted data in Figure 2A). Interestingly, these findings provide evidence of predictive coding of the spatial view cells within the hippocampus. Given that Stachenfeld et al. (2017) previously used the SR model to argue that the hippocampus implements a predictive map, we were interested determining whether the SR model could account for the empirical findings. Stachenfeld et al. (2017) defined the place field as the column of the SR matrix, which allowed them to account for a host of findings in the rodent hippocampus. Here, we use a similar logic to study the firing fields of the spatial view cells in the primate. We found that the SR spatial view cell receptive fields were very similar to the empirical data, thus suggesting that the SR model can account for the predictive coding in monkey spatial view cells (Figure 2).

The SR Model Accounts for Spatial Grid Cells A growing body of influential research has suggested that the primate system contains visual grid cells, which are a specific type of grid cell that fires based on where the animal is looking, rather than where the animal is standing, in the environment (see Introduction). Thus, here we implemented a simulation of the findings from previous work with the SR model, that showed that an eigenvector decomposition of the SR matrix can predict grid-cell-like firing patterns (Stachenfeld et al., 2017). Here, we implemented a graph-world instantiation of tasks from studies that have observed visual grid cells in monkeys (e.g., Killian et al., 2012) and visual-grid-cell-like activity using fMRI in humans (e.g., Julian et al., 2018; Nau, Navarro Schröder, et al., 2018). Specifically, here we created a network of nodes, which represent possible looking locations on the computer screen (Figure 3). Here, we demonstrate that the SR eigenvector visual grid field accounts for the evidence of visual grid coding in the medial entorhinal cortex in both monkeys and humans (Figure 3; $\gamma = 0.95$; we calculated the SR via the analytical method).

Altogether, we provide evidence that the SR model can account both for findings of spatial view cells in the hippocampus and visual grid cells in the medial entorhinal cortex. Therefore, our work here further supports the hypothesis that the hippocampal system encodes a predictive map (Stachenfeld et al., 2017). Moreover, we extend a novel hypothesis that the primate hippocampal system encodes a pre-

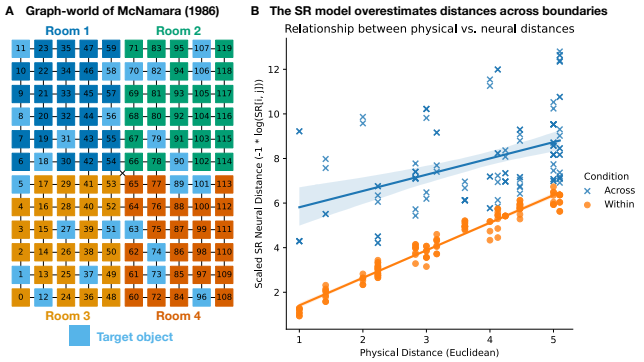


Figure 5: The SR model can account for the finding that participants tend to overestimate distances across a boundary vs. within a boundary. A) We modeled the task from McNamara (1986) using a graph structure. B) Our results indicate that the model overestimated distances across the boundaries vs. within a boundary. Here, we specifically account for the data from McNamara (1986) and we highlight that our findings fit with extant behavioral studies with similar findings.

dictive visually-guided cognitive graph. We now demonstrate how implementing our hypothesis allows us to account for behavioral findings that support the cognitive graph hypothesis.

The SR Model Accounts for Findings Across Several Behavioral Tasks

Given the somewhat contradictory findings and hypotheses of the neuroscientific (e.g., Euclidean/cognitive map) and behavioral (e.g., non-Euclidean/cognitive graph) study of spatial cognition (see Introduction), we were interested in determining whether the SR model could account for behavioral findings that have been used to support the cognitive graph hypotheses. Importantly, the SR model allows us to make quantitative and qualitative predictions based on given (mathematical) assumptions, thus allowing us to more directly test the cognitive graph hypothesis.

The SR Model Accounts for Asymmetric Distance Judgments A body of findings has shown that behavioral judgments of distance can be asymmetric: i.e., AB $\neq$ BA (see Introduction). For example, previous findings have shown that distance judgments can differ based on whether a given location is considered to be a landmark or not (e.g., Burroughs & Sadalla, 1979; Cadwallader, 1979; McNamara, 1991; Mcnamara & Diwadkar, 1997). Thus, the finding of asymmetric distance judgments violates a core assumption of Euclidean geometry that the distance between two locations should be symmetric: i.e., AB = BA. Stachenfeld et al. (2017) showed that place cell coding can be asymmetric in certain conditions (e.g., based on asymmetrical exploration). Thus, we took a similar approach to model behavioral performance.

To determine whether the SR model can account for asymmetric distance judgments, we generated a simple environment, which can be thought of as a form of linear track. We modeled both a symmetric transition structure (Figure 4A) and an asymmetric transition structure (Figure 4C; in both simulations, we used $\gamma = 0.75$ and we calculated the SR using the analytical method). We simulated behavioral distance judgments in the model by indexing the SR matrix using the row index of the first location and the column index of the second location. For plotting, we scaled the resultant value from the SR matrix as follows: $d = -1 \times log(M[i, j])$. The multiplication by -1 scales the data to account such that things that are "further" elicit "further" distance judgments and the *log* transformation deals with the nonlinearity of the SR values. We found that the SR model predicted that simulated distance judgments would be symmetric in the symmetric condition (Figure 4B); however, in the asymmetric condition, the model predicted that distance judgments would be asymmetric (Figure 4D). Therefore, these findings reveal that the SR model can readily account for asymmetric distance judgments that have been reported in the literature.

While our simulation here was intentionally simple, we note that we could easily conceive of methods to scale these simulations up by focusing on the exploration history that participants take in environments. For example, given our focus on the role of vision in guiding the hippocampal cognitive graph, we could imagine that we could use eye tracking data to create training paths for models like the SR model. For example, these findings could provide mechanistic insight into why participants often make asymmetric distance judgments for landmarks. Likewise, we could imagine using brain recording techniques; for example, to determine if participants may reactivate key landmarks as they take paths through the environment.

The SR Model Accounts for Boundary Expansion Effects A classic finding in the behavioral literature is that boundaries tend to cause participants to overestimate distances across boundaries vs. within the same room, which we refer to as the "boundary expansion effect" (see Introduction). McNamara (1986) asked participants learn the locations of objects within 4 rooms. We simulated SR navigation in these rooms on a graph (Figure 5A). We used the same methods as our simulation of the asymmetric distances (see previous simulation; $\gamma = 0.7$; we calculated the SR via the analytical method). We restricted the comparison of between-room nodes to only include pairs of "target" locations that were within the same distance distribution as the within-room distance distribution (i.e., to match the overall distances between conditions). We found that the distance judgments in the "across boundary" condition were overestimated relative to the "within boundary" condition (i.e., notice the points for the "across" condition shift up relative to the "within" condition in Figure 5B), thus accounting for the "boundary expansion" effect. In summary, we demonstrate that the SR model can readily account for the boundary expansion effect, which has been observed in numerous studies (Cohen et al., 1978; Kosslyn et al., 1974; McNamara, 1986; Newcombe & Liben, 1982; also see: Hirtle & Jonides, 1985; McNamara et al., 1989), and we highlight

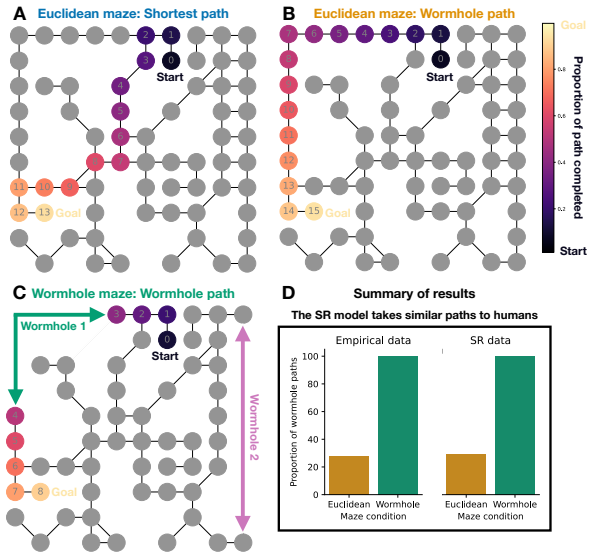


Figure 6: The SR model takes similar paths to humans in both a Euclidean and a wormhole environment. Specifically, here, we created a graph-world simulation of the task created by Warren et al. (2017). The model takes the shortest path the majority of the time for both the "Euclidean maze" (i.e., the path shown in panel A), only taking the alternative "wormhole 1" path 29% of the time (human participants took this path 28% of the time; panels B and D). For the wormhole maze, both the SR model and human participants took the "wormhole 1" path 100% of the time (panels C and D).

that these findings could be leveraged to account for a variety of measures, including distance judgments and priming.

The SR Model Accounts for Human Navigation within Euclidean and Non-Euclidean Mazes Perhaps some of the strongest evidence for the cognitive graph hypothesis comes from tasks in which participants navigate in non-Euclidean mazes (see Introduction). For example, Warren et al. (2017) created two versions of a maze environment: one with a typical, Euclidean layout and another with wormholes that teleported the participants over long distances and also angular changes in the environment (i.e., a non-Euclidean). Note that the wormhole maze is physically impossible within a 2D layout (i.e., it is non-Euclidean). We generated graph-world environments for simulating performance in the SR ($\gamma = 0.9$; we calculated the SR via the temporal difference rule with 1,350 steps in a weighted random walk to simulate human exploration; 100 simulated participants).

In the Euclidean maze, the SR model and humans both took the shortest path between target nodes the majority of the time (Figure 6A), rarely taking the "wormhole" paths (28% of humans took the wormhole paths in this condition; similarly, the SR model took the "wormhole 1" path 29% of the time in this condition; Figure 6B and D). In the wormhole condition, both the SR model and human participants took the "wormhole 1" path 100% of the time (Figure 6C and D). Moreover, when Warren et al. (2017) asked participants about the mazes in a post-experiment questionnaire, none of the participants in the wormhole mazed noticed that anything was amiss about the environment. While Warren et al. (2017) concluded that their results provide evidence for the cognitive graph hypothesis, our results provide much more direct support given that the SR model was only given access to learning an underlying graph of the environments (i.e., our results provide qualitative and quantitative support for their verbal theories).

The SR Model Accounts for Findings that Spatial Memory Resolution Differs by the Availability of Visual Cues In an influential review, Peer et al. (2021) proposed that humans can make use of both a cognitive-map system and a cognitive-graph system depending on the nature of the environment and the task demands. Thus, they argue that there might not be a universal, underlying representation for space, rather the representations that are brought to bear may depend on a number of experiment-specific factors. Thus, their framework can account for the disparate findings in the literature that have been used to suggest that spatial representations are more Euclidean (e.g., the cognitive map hypothesis) vs. more non-Euclidean and graph-like (e.g., the cognitive graph hypothesis). Peer, Nadar, and Epstein (2024) followed up on this hypothesis by designing mazes with varying constraints on the participants movements through the space, thus allowing them to test if participants form more map-like vs. more graph-like representations of the space. They found that, across a variety of measures, the resolution of spatial memory varied as a function of the visual cues that were available to participants; thus, they argued that they supported their theoretical framework by showing more map-like representations being favored in open environments and more graph-like representations being favored in environments with boundaries.

We were interested in discovering whether we could capture the findings from Peer et al. (2024) within the SR model, which would perhaps suggest we should reconceptualize their findings as being equally (if not more) supported by an underlying cognitive graph that varies as a function of its connectedness rather than as reflecting a fundamentally different type of representation. Thus, we created 3 graph-world environments to simulate their experimental conditions (see Figure 7 A–C). The novel advance of our approach here is that we argue that humans (and other primates) can use vision to "visually" traverse the environment without moving. We believe that our assumption is supported by the data that we described in our simulation of spatial view cells and spatial grid cells (see above), among other previously published findings and interpretations, especially of spatial view cells. Therefore, in our simulations here, we model the "open" environment as a relatively complete grid of possible locations within the broader environment because humans would have been able to make saccades throughout the environment from any given position. In contrast, visual exploration would be limited to

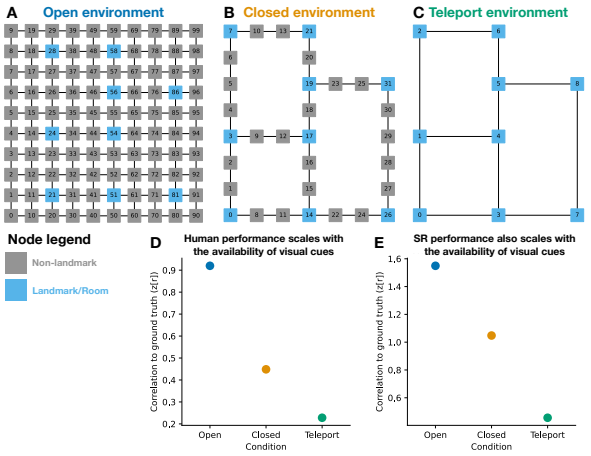


Figure 7: The SR model can account for the finding that the apparent resolution of human spatial memory scales with the availability of visual cues. We modeled the task from Peer et al. (2024), with 3 graph-world environments: (A) "open": all landmarks visible from any location, (B) "closed": visual exploration was limited to corridors, and (C) "teleport": no visual exploration between landmarks. Humans and the SR model performed best on the "open"m followed by the "closed", and worst on the "teleport" condition (compare panels D and E). Abbreviation: z[r] = Fisher's r-to-z transformation for the Pearson correlation coefficient.

corridors in the "closed" environment and completely eliminated in the "teleport" environment ($\gamma = 0.99$, we calculated the SR via the analytical method).

Here, for the sake of brevity and for the nature of the standard version of the SR model, we focus our attention on the distance-estimation task. We simulated distance judgments by correlating the values from the SR matrix with the ground-truth Euclidean distances[1]. We note that both SR and human performance differed as a function of the environment, such that the correlation between the ground-truth and distance judgments was highest for the "open" condition, followed by the "closed" condition, and worst on the "teleport" condition. Thus, although Peer et al. (2024) interpreted their findings to indicate a qualitative difference in the representations that participants use to solve the tasks (i.e., a difference in the reliance on cognitive maps vs. cognitive graphs), our results suggest that we can instead reconceptualize their findings as differing as a measure of quantitative differences (e.g., a different number of nodes) rather than qualitative difference (e.g., all tasks can be solved by an underlying cognitive graph). We would also like to generalize our claims here to a broader discussion: previous results that support the notion of a cognitive map map may just be relatively well-connected cognitive graphs masquerading as being more metric, Euclidean, globally consistent representations.

[1]We multiplied the correlation values by -1 (i.e. since the SR values are an inverse proxy of distance).

Conclusion

Our results provide an important advancement to theories of the primate hippocampus and spatial memory. We show that by modeling the influence of vision, we can account for a range of neuroscientific findings that then support our use of these assumptions in modeling behavioral performance. Importantly, our findings show that our novel use of the SR model can readily account for extant behavioral findings in the literature that have been used to support the cognitive graph hypothesis, including asymmetric distance judgments, boundary expansion effects, taking novel paths in wormhole environments, and the role that visual exploration plays on shaping the apparent resolution of spatial memory. We would also like to highlight our novel hypothesis that some of the previous literature that has been used to support the cognitive map/Euclidean hypothesis may be influenced by highly connected graphs masquerading as map-like knowledge (e.g., due to the "vista" space nature of the environments that allow for rapid visual exploration). We would also like to emphasize that other recent work with the SR model has shown that under conditions in which there are boundaries and non-rectangular environments, that the modeled representations of hippocampal activity and grid-cell-like activity are warped (e.g., Stachenfeld et al., 2017; Bellmund et al., 2019). Interestingly, Bellmund et al. (2019) found that the predicted deformations in the SR eigenvector grid fields correlated with changes in human behavioral performance.

We also note that our overall computational approach is not entirely different from recent work modeling spatial view cells (e.g., De Araujo, Rolls, & Stringer, 2001; Rolls, 2023; Stringer, Rolls, & Trappenberg, 2005). Indeed, the main difference in their approach is they used attractor neural networks to model neuroscientific findings rather than behavior. We agree with other theoretical accounts that the computational architecture of the hippocampus may make it especially well suited for computing the SR (e.g., perhaps the CA3 fields; cf. Rolls, 2023; Stachenfeld et al., 2017). Thus, we think that future research can seek to further explore the ideas we lay out here as well as how they compare with other research. Likewise future research can adjudicate between more neuroscientifically guided versions of the SR model (e.g., modeling things like place cells; De Cothi & Barry, 2020) vs. the graph-based version we used here. Such explorations can further adjudicate the conditions in which graph-based vs. more metric-based information better account for spatial memory. Based on our results here, we argue that a rich set of data can be accounted for within a single graph-based network that varies as a function of connectedness, rather than as fundamentally different representations.

In summary, we propose a novel hypothesis that primate spatial cognition is supported by a visually-guided cognitive graph. We provide initial support for our hypothesis, and we will continue to further advance our framework within novel computational architectures, including additional modeling of results from other papers and generating novel predictions.

Acknowledgements

We thank the Department of Psychology at Colby College for supporting this work. We also thank other members of our lab for helpful conversations.

References

Bellmund, J. L. S., de Cothi, W., Ruiter, T. A., Nau, M., Barry, C., & Doeller, C. F. (2019). Deforming the metric of cognitive maps distorts memory. *Nature Human Behaviour*, *4*(2), 177–188. doi: 10.1038/s41562-019-0767-3

Bellmund, J. L. S., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*, *362*(6415), eaat6766. doi: 10.1126/science.aat6766

Burroughs, W. J., & Sadalla, E. K. (1979). Asymmetries in Distance Cognition. *Geographical Analysis*, *11*(4), 414–421. doi: 10.1111/j.1538-4632.1979.tb00709.x

Byrne, P., Becker, S., & Burgess, N. (2007). Remembering the past and imagining the future: A neural model of spatial memory and imagery. *Psychological Review*, *114*(2), 340–375. doi: 10.1037/0033-295X.114.2.340

Cadwallader, M. (1979). Problems in Cognitive Distance: Implications for Cognitive Mapping. *Environment and Behavior*, *11*(4), 559–576. doi: 10.1177/0013916579114007

Chrastil, E. R., & Warren, W. H. (2014). From cognitive maps to cognitive graphs. *PloS One*, *9*(11), e112544. doi: 10.1371/journal.pone.0112544

Cohen, R., Baldwin, L. M., & Sherman, R. C. (1978). Cognitive Maps of a Naturalistic Setting. *Child Development*, *49*(4), 1216. doi: 10.2307/1128763

De Araujo, I., Rolls, E., & Stringer, S. (2001). A view model which accounts for the spatial fields of hippocampal primate spatial view cells and rat place cells. *Hippocampus*, *11*(6), 699–706. doi: 10.1002/hipo.1085

De Cothi, W., & Barry, C. (2020). Neurobiological successor features for spatial navigation. *Hippocampus*, *30*(12), 1347–1355. doi: 10.1002/hipo.23246

Ekstrom, A. D. (2015). Why vision is important to how we navigate: Human Spatial Navigation and Vision. *Hippocampus*, *25*(6), 731–735. doi: 10.1002/hipo.22449

Gallistel, C. R. (1990). *The organization of learning*. Cambridge (Mass.) London: the MIT press.

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–806. doi: 10.1038/nature03721

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

Julian, J. B., Keinath, A. T., Frazzetta, G., & Epstein, R. A. (2018). Human entorhinal cortex represents visual space using a boundary-anchored grid. *Nature Neuroscience*, *21*(2), 191–194. doi: 10.1038/s41593-017-0049-1

Killian, N. J., Jutras, M. J., & Buffalo, E. A. (2012). A map of visual space in the primate entorhinal cortex. *Nature*, *491*(7426), 761–764. doi: 10.1038/nature11587

Killian, N. J., Potter, S. M., & Buffalo, E. A. (2015). Saccade direction encoding in the primate entorhinal cortex during visual exploration. *Proceedings of the National Academy of Sciences*, *112*(51), 15743–15748. doi: 10.1073/pnas.1417059112

Kosslyn, S. M., Pick, H. L., & Fariello, G. R. (1974). Cognitive Maps in Children and Men. *Child Development*, *45*(3), 707. doi: 10.2307/1127837

McNamara, T. P. (1986). Mental representations of spatial relations. *Cognitive Psychology*, *18*(1), 87–121. doi: 10.1016/0010-0285(86)90016-2

McNamara, T. P. (1991). Memory's View of Space. In *Psychology of Learning and Motivation* (Vol. 27, pp. 147–186). Elsevier. doi: 10.1016/S0079-7421(08)60123-1

Mcnamara, T. P., & Diwadkar, V. A. (1997). Symmetry and Asymmetry of Human Spatial Memory. *Cognitive Psychology*, *34*(2), 160–190. doi: 10.1006/cogp.1997.0669

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), 211–227. doi: 10.1037/0278-7393.15.2.211

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), 663–678. doi: 10.1038/nrn1932

Moser, E. I., & Moser, M.-B. (2008). A metric for space. *Hippocampus*, *18*(12), 1142–1156. doi: 10.1002/hipo.20483

Nau, M., Julian, J. B., & Doeller, C. F. (2018). How the Brain's Navigation System Shapes Our Visual Experience. *Trends in Cognitive Sciences*, *22*(9), 810–825. doi: 10.1016/j.tics.2018.06.008

Nau, M., Navarro Schröder, T., Bellmund, J. L. S., & Doeller, C. F. (2018). Hexadirectional coding of visual space in human entorhinal cortex. *Nature Neuroscience*, *21*(2), 188–190. doi: 10.1038/s41593-017-0050-8

Newcombe, N., & Liben, L. S. (1982). Barrier effects in the cognitive maps of children and adults. *Journal of Experimental Child Psychology*, *34*(1), 46–58. doi: 10.1016/0022-0965(82)90030-3

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. doi: 10.1016/0006-8993(71)90358-1

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

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. doi: 10.1016/j.tics.2020.10.004

Peer, M., Nadar, C., & Epstein, R. A. (2024). The format of the cognitive map depends on the structure of the envi-

ronment. *Journal of Experimental Psychology: General*, *153*(1), 224–240. doi: 10.1037/xge0001498

Rolls, E. T. (2023). Hippocampal spatial view cells, place cells, and concept cells: View representations. *Hippocampus*, *33*(5), 667–687. doi: 10.1002/hipo.23536

Rolls, E. T., & O'Mara, S. M. (1995). View-responsive neurons in the primate hippocampal complex. *Hippocampus*, *5*(5), 409–424. doi: 10.1002/hipo.450050504

Rolls, E. T., Robertson, R. G., & Georges-François, P. (1997). Spatial View Cells in the Primate Hippocampus. *European Journal of Neuroscience*, *9*(8), 1789–1794. doi: 10.1111/j.1460-9568.1997.tb01538.x

Russek, E. M., Momennejad, I., Botvinick, M. M., Gershman, S. J., & Daw, N. D. (2017). Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLOS Computational Biology*, *13*(9), e1005768. doi: 10.1371/journal.pcbi.1005768

Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, *20*(11), 1643–1653. doi: 10.1038/nn.4650

Stevens, A., & Coupe, P. (1978). Distortions in judged spatial relations. *Cognitive Psychology*, *10*(4), 422–437. doi: 10.1016/0010-0285(78)90006-3

Stringer, S., Rolls, E., & Trappenberg, T. (2005). Self-organizing continuous attractor network models of hippocampal spatial view cells. *Neurobiology of Learning and Memory*, *83*(1), 79–92. doi: 10.1016/j.nlm.2004.08.003

Taube, J. S., Muller, R., & Ranck, J. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *The Journal of Neuroscience*, *10*(2), 420–435. doi: 10.1523/JNEUROSCI.10-02-00420.1990

Tversky, B. (1992). Distortions in cognitive maps. *Geoforum*, *23*(2), 131–138. doi: 10.1016/0016-7185(92)90011-R

Wallace, D. J., Greenberg, D. S., Sawinski, J., Rulla, S., Notaro, G., & Kerr, J. N. D. (2013). Rats maintain an overhead binocular field at the expense of constant fusion. *Nature*, *498*(7452), 65–69. doi: 10.1038/nature12153

Warren, W. H. (2019). Non-Euclidean navigation. *The Journal of Experimental Biology*, *222*(Suppl 1), jeb187971. doi: 10.1242/jeb.187971

Warren, W. H., Rothman, D. B., Schnapp, B. H., & Ericson, J. D. (2017). Wormholes in virtual space: From cognitive maps to cognitive graphs. *Cognition*, *166*, 152–163. doi: 10.1016/j.cognition.2017.05.020

Wirth, S., Baraduc, P., Planté, A., Pinède, S., & Duhamel, J.-R. (2017). Gaze-informed, task-situated representation of space in primate hippocampus during virtual navigation. *PLOS Biology*, *15*(2), e2001045. doi: 10.1371/journal.pbio.2001045