Individual Variation in Non-Spatial Navigation is Related to Adaptive Learning and Use of Successor Representations

Alexa K. Booras [1] **Theodoros Kapogianis** [2] **Aaron M. Bornstein**[1,*] **Elizabeth R. Chrastil**[1,2,*]

[1] Department of Cognitive Sciences, University of California, Irvine
[2] Department of Neurobiology and Behavior, University of California, Irvine

`boorasa@uci.edu, tkapogia@uci.edu, aaron.bornstein@uci.edu, chrastil@uci.edu`

Abstract

A crucial component of decision-making is the ability to learn from experience and form internal representations, or 'cognitive graphs.' Here we relate participant learning of different non-spatial associative networks to their subsequent inferences about the latent graph. Participants first actively explored and implicitly learned stimulus-stimulus associations by viewing groups of images (nodes) representing potential paths (edges) and selecting images which triggered the presentation of the next image group. Next, participants were tested on their ability to navigate quickly through the graph. We analyzed how participant exploration related to their later navigation performance using a *successor representation* (SR) fit to exploration response times, capturing individual differences in planning horizon as a function of graph structure. We found that people who changed their planning horizon to better align with the optimal planning horizon of a particular graph had better knowledge of that graph. Additionally, we found that individual differences in memory specificity constrained the ability to adjust to the least-structured graph (longest planning horizon). Overall, our results suggest that humans adapt their learning and use of internal representations during associative network learning in ways that reflect their environmental regularities and memory constraints.

Background

How do people build memories that help them later make decisions? The regularities of our environment shape our memory for events, and our memory of those regularities are critical to planning and navigation in later encounters (Noh* et al., 2026; Peer et al., 2024; Rmus et al., 2022; Yoo et al., 2024). Such *latent-learned* associative representations can usefully be understood as cognitive graphs: "edges" within these graphs reflecting associations between "nodes" (Chrastil and Warren, 2014). Organizing information into cognitive graphs that reflect underlying environmental regularities and structure may aid in learning and memory (Coutrot et al., 2022; Kahn et al., 2018; Karuza et al., 2019; Lynn and Bassett, 2020; Schapiro et al., 2013).

One biologically plausible mechanism (Lee, 2022) by which people learn compact representations of cognitive graphs is the *successor representation* (SR), which encodes each state of the environment in terms of the discounted future occupancy of other states (Dayan, 1993; Stachenfeld et al., 2017). An SR can model individuals' navigation decisions in a community-structured state space, and when fit to response times better captures performance on a non-spatial navigation task than comparison models (Wientjes and Holroyd, 2024). Despite SR's ubiquity in research (Boorman et al., 2021; Gershman et al., 2012; Son et al., 2023; Tacikowski et al., 2024), it remains unknown whether people can adapt their internal planning horizon - how far in the future predictive associations extend - to best fit the structure of their environment.

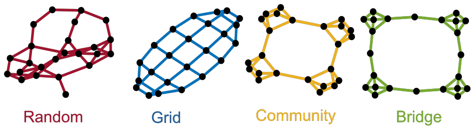

Figure 1: Graph structures.

In this experiment, individuals explored different graph-structured state spaces (Fig. 2, 1). We predicted that a successor representation would model individuals' navigation responses and that individuals would be able to adapt their planning horizon to the specific pattern of long- and short-range regularities in each type of graph. Indeed, we found that individuals who can flexibly adapt their internal representations to best suit their environments have better navigation performance across various non-spatial graphs. Furthermore, in an unstructured graph, this correlated to the individual's memory specificity.

Methods

Study Design We recruited 58 male and 58 female young and healthy adults (18-35 years old) through Prolific. Each participant completed a learning phase (4 minutes), navigation test phase (10 trials), and distance judgment phase (48 trials) for four different graph structures (Fig. 1). See Figure 2 for task description. On day two of the experiment, participants underwent the Mnemonic Similarity Task (MST) - an established object

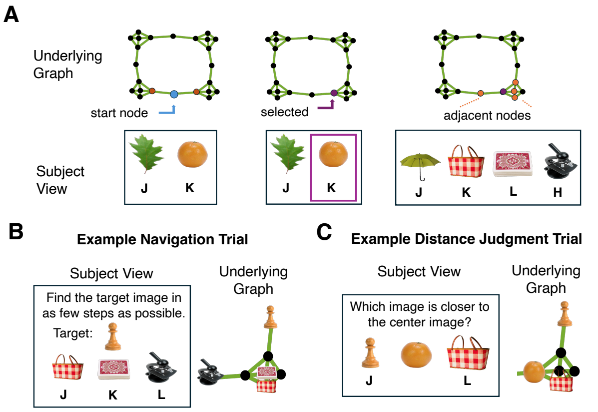

Figure 2: A) Example exploration trial. Participants were shown groups of images. When they selected an image, the next group presented was adjacent on the underlying graph. Participants were never shown the underlying graph. B) Example navigation test trial. Participants were tasked with finding the target node within 20 seconds. C) Example distance judgment test trial. Participants viewed three images and were asked whether the left or right image was closer to the center image.

recognition task designed to test one's ability to remember and distinguish between similar items (Stark et al., 2015, 2019).

Successor Representation We modeled participants' knowledge of the graphs as a successor representation (SR). The agent learns a matrix M where there are s states in the environment. Learning is calculated using the temporal difference algorithm (Dayan, 1993). γ indicates the planning horizon (discount factor) and λ indicates the learning rate, fixed at 0.1 (Wientjes and Holroyd, 2024, Pudhiyidath et al., 2022).

$$M[s_t,:] = M[s_t,:] + \lambda(1_{s_t} + \gamma M[s_{t+1},:] - M[s_t,:])\ . \tag{1}$$

The planning horizon parameter indicates how far into the future predictions extend. To determine the optimal planning horizon for each graph, we simulated random walks along the graph at a range of planning horizons (.01 to .99 with a step size of .01) and calculated shortest path distances. The shortest possible path of the actual graph was then subtracted from the shortest possible path of the successor representation, averaged across start/target pairs.

We fit the SR to participants' trial-by-trial exploration response times in each of the four graphs, modeling RTs as a regression on the mean parameter of a shifted lognormal distribution. Our model and priors were taken from Wientjes and Holroyd, 2024:

$$y_{ij} - ndt_i \sim lognormal(\alpha + \beta x_i, \theta_i) \tag{2}$$

Where i represents the participant and j represents the trial.

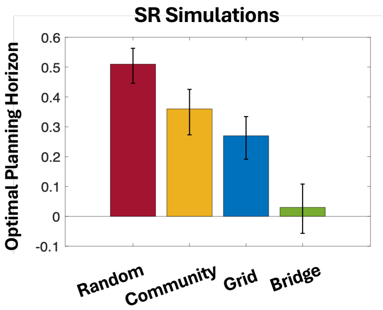

Figure 3: 10,000 Random walks for each graph were simulated and the shortest possible path calculated across different planning horizons.

Results

We expected participants would adapt their planning horizon to each graph. Our simulations revealed the bridge graph to have the shortest planning horizon, therefore we used each individuals' bridge graph γ as their baseline, subtracting it from those fit to the other graphs. We found that individuals whose planning horizons adapted to a greater degree between random and bridge also had fewer excess steps ($r = -.28$, $p = .006$), and similarly for community and bridge ($r = -.30$, $p = .004$) and grid and bridge ($r = -.23$, $p = .03$). Memory specificity (Lure Discrimination Index) positively correlated to planning horizon on the random graph ($r = .28$, $p = .005$, all other $r < .15$, $p > .12$).

Discussion

The planning horizon is a key individual difference measuring how far into the future one extends their predictions. Here we expand on previous work indicating the importance of predictive representations for navigating structured environments (Dayan, 1993; Stachenfeld et al., 2017; Wientjes and Holroyd, 2024). We found individual differences in planning horizons that relate to improvements in graph learning. The more people adapted their planning horizon between graph structures to be better aligned with the optimal planning horizon, the more efficient their navigation. Intuitively, in the bridge graph, a short planning horizon may be beneficial because it allows for stronger learning of local associations, while in the random graph a longer planning horizon allows for learning varied associations. However, a longer planning horizon may require stronger memory abilities. This is supported by our results correlating better memory specificity to longer planning horizon on the random graph. Overall, our results suggest that representational flexibility during exploration benefits learning of non-spatial graph structures, and ones representations may be constrained by memory in the absence of structural regularities.

Acknowledgments

The authors acknowledge funding from NIMH T32MH119049 and NINDS R01NS11946.

References

Boorman, E. D., Sweigart, S. C., & Park, S. A. (2021). Cognitive maps and novel inferences: A flexibility hierarchy. *Current Opinion in Behavioral Sciences*, *38*, 141–149.

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

Coutrot, A., Manley, E., Goodroe, S., Gahnstrom, C., Filomena, G., Yesiltepe, D., Dalton, R. C., Wiener, J. M., Hölscher, C., Hornberger, M., et al. (2022). Entropy of city street networks linked to future spatial navigation ability. *Nature*, *604*(7904), 104–110.

Dayan, P. (1993). Improving generalization for temporal difference learning: The successor representation. *Neural computation*, *5*(4), 613–624.

Gershman, S. J., Moore, C. D., Todd, M. T., Norman, K. A., & Sederberg, P. B. (2012). The successor representation and temporal context. *Neural Computation*, *24*(6), 1553–1568.

Kahn, A. E., Karuza, E. A., Vettel, J. M., & Bassett, D. S. (2018). Network constraints on learnability of probabilistic motor sequences. *Nature human behaviour*, *2*(12), 936–947.

Karuza, E. A., Kahn, A. E., & Bassett, D. S. (2019). Human sensitivity to community structure is robust to topological variation. *Complexity*, *2019*(1), 8379321.

Lee, H. (2022). Toward the biological model of the hippocampus as the successor representation agent. *Biosystems*, *213*, 104612.

Lynn, C. W., & Bassett, D. S. (2020). How humans learn and represent networks. *Proceedings of the National Academy of Sciences*, *117*(47), 29407–29415.

Noh*, S. M., Cooper*, K. W., Guo, S., Zhou, D., Stark, C. E., & Bornstein, A. M. (2026). Multi-step inference can be improved across the lifespan with individualized memory interventions. *The Journals of Gerontology, Series B: Psychological Sciences and Social Sciences*.

Peer, M., Nadar, C., & Epstein, R. A. (2024). The format of the cognitive map depends on the structure of the environment. *Journal of Experimental Psychology: General*, *153*(1), 224.

Pudhiyidath, A., Morton, N. W., Viveros Duran, R., Schapiro, A. C., Momennejad, I., Hinojosa-Rowland, D. M., Molitor, R. J., & Preston, A. R. (2022). Representations of temporal community structure in hippocampus and precuneus predict inductive reasoning decisions. *Journal of Cognitive Neuroscience*, *34*(10), 1736–1760.

Rmus, M., Ritz, H., Hunter, L. E., Bornstein, A. M., & Shenhav, A. (2022). Humans can navigate complex graph structures acquired during latent learning. *Cognition*, *225*, 105103.

Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature neuroscience*, *16*(4), 486–492.

Son, J.-Y., Bhandari, A., & FeldmanHall, O. (2023). Abstract cognitive maps of social network structure aid adaptive inference. *Proceedings of the National Academy of Sciences*, *120*(47), e2310801120.

Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature neuroscience*, *20*(11), 1643–1653.

Stark, S. M., Kirwan, C. B., & Stark, C. E. (2019). Mnemonic similarity task: A tool for assessing hippocampal integrity. *Trends in cognitive sciences*, *23*(11), 938–951.

Stark, S. M., Stevenson, R., Wu, C., Rutledge, S., & Stark, C. E. (2015). Stability of age-related deficits in the mnemonic similarity task across task variations. *Behavioral neuroscience*, *129*(3), 257.

Tacikowski, P., Kalender, G., Ciliberti, D., & Fried, I. (2024). Human hippocampal and entorhinal neurons encode the temporal structure of experience. *Nature*, *635*(8037), 160–167.

Wientjes, S., & Holroyd, C. B. (2024). The successor representation subserves hierarchical abstraction for goal-directed behavior. *PLOS Computational Biology*, *20*(2), e1011312.

Yoo, J., Chrastil, E. R., & Bornstein, A. M. (2024). Cognitive graphs: Representational substrates for planning. *Decision*, *11*(4), 537–556.