Replay shapes abstract cognitive maps for efficient social navigation

Jae-Young Son (jae@brown.edu)
Dept. of Cognitive, Linguistic, and Psychological Sciences, Brown University

Marc-Lluís Vives (m.l.vives.moya@fsw.leidenuniv.nl)
Institute of Psychology, Leiden University

Apoorva Bhandari (apoorva_bhandari@brown.edu)
Dept. of Cognitive, Linguistic, and Psychological Sciences, Brown University

Oriel FeldmanHall (oriel.feldmanhall@brown.edu)
Dept. of Cognitive, Linguistic, and Psychological Sciences, Brown University
Carney Institute for Brain Science, Brown University

Abstract:

To make adaptive social decisions, people must anticipate how information flows through their social network. While this requires knowledge of how people are connected, networks are too large to have firsthand experience with every possible route between individuals. How, then, are people able to accurately track information flow through social networks? We find that people cache abstract knowledge about social network structure as they learn who is friends with whom, which enables the identification of efficient routes between remotely-connected individuals. These cognitive maps of social networks, which are built immediately after learning, are then reshaped through overnight rest. During these extended periods of rest, a replay-like mechanism helps to make these maps increasingly abstract, which especially privileges improvements in social navigation accuracy for the longest communication paths spanning distinct communities. Together, these findings provide mechanistic insight into the sophisticated mental representations humans use for social navigation.

Keywords: social networks; cognitive maps; successor representation; decision making

Introduction

Human social life is embedded within a complex web of connections. Navigating through the social world therefore requires representing people's relationships with one another, including those extending beyond one's direct circle of friends (Basyouni & Parkinson, 2022). However, little is known about the mental representations that enable social navigation through complex networks, nor how these representations are built from limited direct experience.

Recent work points to multistep abstraction as a promising mechanism for representing cognitive maps of social networks, as it encodes not only the direct connections between entities (e.g., friendships), but also longer-range, multistep connections like friends-of-friends (Son et al., 2023). Learning abstract cognitive maps affords rapid inference about distant relations, which likely aids social navigation.

Although abstract representations can be learned 'online' from direct experience, past work shows that the brain can efficiently reuse experiences during 'offline' replay to simulate synthetic sequences, which aids learning of more abstract representations (Momennejad, 2020). It is therefore likely that offline replay contributes to building the kinds of abstract representations needed for longer-range social navigation through social networks.

In this work, we test whether humans rely on cognitive maps of social networks for social navigation and whether a replay-like mechanism supports more successful social navigation following overnight rest.

Methods

Overview. To assess how humans solve social navigation problems, we created a task where subjects learn about friendships in an artificial social network, allowing us to probe whether subjects could navigate information flow through the network. We had subjects take the navigation task immediately after learning, then brought subjects back to test whether navigation accuracy improved after overnight rest. Using computational modeling, we tested whether subjects show evidence of using an abstract cognitive map, and whether a replay-like mechanism helps to scaffold more successful social navigation.

Learning task. Subjects learned about friendships from a 'flashcard' game (Figure 1A). On each trial, subjects were shown a 'Target' network member, and were required to find all of the Target's friends amongst the remaining face-down cards. Cards flipped face-up when subjects made correct responses; cards remained face-down for incorrect responses.

Social navigation task. On each trial, subjects were told a network member needed to choose between Sources A and B for passing a letter to a particular Target within the network (e.g., if Source A were chosen, A would pass the letter to one of their friends, who would pass it to one of *their* friends, and so on until the letter reached the Target; Figure 1B). The subject's task was to choose the Source that would result in the most efficient delivery. An accurate response was defined as choosing the Source with the shortest path to the Target. The Target changed on every trial, such that successful navigation required flexible use of knowledge about connections between network members. No feedback was ever provided.

Computational model of replay. We modeled abstraction using the Successor Representation (SR), which approximates the probability of transitioning from a Source to a Target in a given number of steps, $p(T \mid S, \gamma)$ (Dayan, 1993; Russek et al., 2017). The parameter γ controls how many steps are integrated over, and therefore how abstract the representation is. As $\gamma \rightarrow 0$, the agent represents shorter-range relations, such that the SR only encodes one-step relations (i.e., direct friendships) when $\gamma=0$. As $\gamma \rightarrow 1$, the agent learns longer-range connections (e.g., friends-of-friends). In our implementation, the representation is learned using standard delta-rule updating.

In the SR, replay is a natural mechanism for explaining how overnight rest improves social navigation (Momennejad, 2020). The knowledge cached by the SR is sensitive to an agent's observations, which could include either direct experience from the environment or synthetic experience from offline replay. We hypothesized that an agent's ability to successfully solve longer-range navigation problems depends on building increasingly abstract representations integrating over a greater number of multistep relations (i.e., with larger γ). Intuitively, replay sequences are likely to be shorter during awake rest than during sleep, and it is therefore possible that overnight rest helps to stitch knowledge of pairwise relationships into longer sequences of multistep relations, allowing an agent to build more abstract cognitive maps.

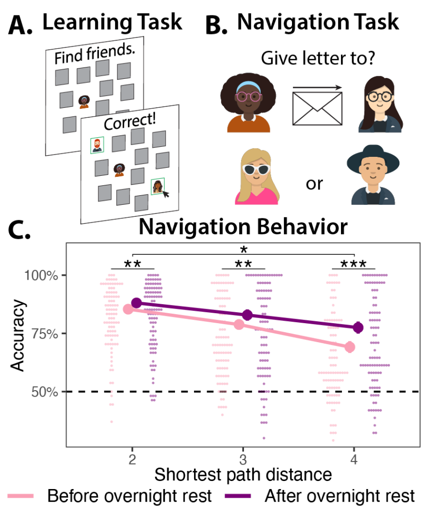

Figure 1. A. Learning task. **B.** Social navigation task. **C.** Navigation accuracy before and after overnight rest.

Results & Discussion

Humans can accurately navigate social networks immediately after learning. Subjects (*N* = 146) achieved above-chance navigation accuracy not only for problems where the Source was directly friends with the Target ('distance-2' accuracy = 80%, *β* = 1.68, 95% CI = [1.45, 1.91], *p* < .001), but also for longer-range problems involving friends-of-friends (distance-3 accuracy = 70%, *β* = 1.06, 95% CI = [0.84, 1.28], *p* < .001), as well as friends-of-friends-of-friends (distance-4 accuracy = 63%, *β* = 0.66, 95% CI = [0.44, 0.87], *p* < .001; Figure 1C). These results suggest humans can learn a cognitive map supporting flexible, long-range social navigation, even when they are only able to observe pairwise friendships in a social network.

Multistep abstraction is sufficient for accurate navigation. Simulation results reveal high navigation accuracy can be explained by multistep abstraction: higher values of γ are associated with greater accuracy for longer-range problems (Figure 2A).

Social navigation improves after rest. After rest, subjects' (*N* = 96) social navigation improved across all distances (distance-2 accuracy = 82%, *β* = 0.23, 95% CI = [0.08, 0.39], *p* = .004; distance-3 accuracy = 75%, *β* = 0.26, 95% CI = [0.09, 0.44], *p* = .003; distance-4 accuracy = 71%, *β* = 0.43, 95% CI = [0.27, 0.59], *p* < .001; Figure 1C). This improvement was particularly pronounced for the longest-range distance-4 problems compared to distance-2 problems (*β* = 0.20, 95% CI = [0.03, 0.36], *p* = .018).

Replay aids social navigation by enabling representation of more abstract cognitive maps. Computational modeling results reveal a significant increase in γ after rest (Day 1 median γ = 0.51, Day 2 median γ = 0.66, one-tailed *p* = .023; Figure 2B). To verify that increases in γ are associated with greater navigation accuracy, we used Spearman rank correlation to test whether changes in γ track changes in accuracy for shorter- and longer-range navigation problems. Indeed, increased γ on Day 2 was associated with improved navigation accuracy for the longer-range problems (distance-3 ρ = 0.21, one-tailed *p* = .022; distance-4 ρ = 0.45, one-tailed *p* < .001; Figure 2C), but not for the shorter-range problems (distance-2 ρ = -0.18, one-tailed *p* = .960).

Together, our results point to multistep abstraction as a representational format enabling social information flow to be tracked in the human mind, and to the importance of a replay-like mechanism for building more abstract cognitive maps supporting longer-range navigation through a social network.

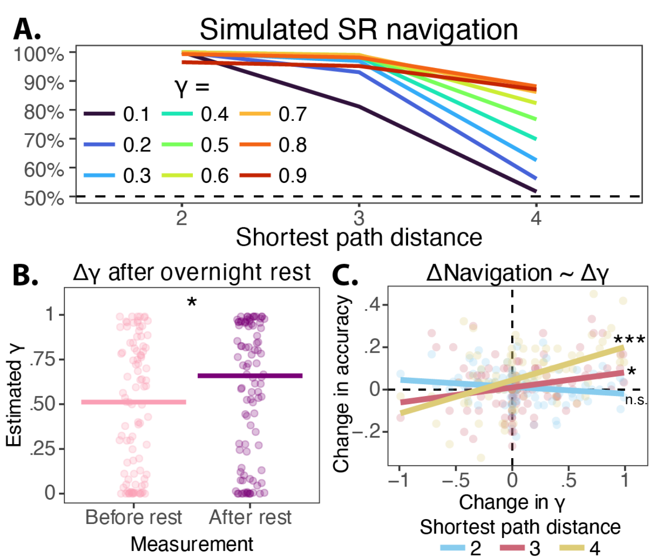

Figure 2. A. Simulated behavior from a computational model of multistep abstraction. **B.** Increased abstraction after rest. **C.** Increased abstraction is associated with better longer-range navigation after rest.

Acknowledgments

We thank the following people for assisting with data collection: Isabella Aslarus, Kayleigh Danowski, Elizabeth Duchan, Yi-Fei Jerry Hu, Alexus Lawrence, Jonathan Palfy, Vera Poyraz, Mehak Malhotra, Maya Mazumder, Samantha Shulman, Ariel Stein, Sofía Vaca Narvaja, and Jenny Wang. We thank Armin Maddah for developing some of the task code used in these studies. Part of this research was conducted using computational resources and services at the Center for Computation and Visualization, Brown University. Advanced access to these computing resources was supported by NIH award 1S10OD025181. This work is supported by the National Science Foundation award 2123469 (O.F.H. and A.B.).

References

Basyouni, R., & Parkinson, C. (2022). Mapping the social landscape: tracking patterns of interpersonal relationships. *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2021.12.006

Dayan, P. (1993). Improving Generalization for Temporal Difference Learning: The Successor Representation. *Neural Computation*, *5*(4), 613-624. https://doi.org/10.1162/neco.1993.5.4.613

Momennejad, I. (2020). Learning Structures: Predictive Representations, Replay, and Generalization. *Current Opinion in Behavioral Sciences*, *32*, 155-166. https://doi.org/https://doi.org/10.1016/j.cobeha.2020.02.017

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. https://doi.org/10.1371/journal.pcbi.1005768

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. https://doi.org/10.1073/pnas.2310801120