Check for updates

ASSOCIATION FOR PSYCHOLOGICAL SCIENCE

Current Directions in Psychological Science
1–8

DOI: 10.1177/09637214251342742
www.psychologicalscience.org/CDPS
Sage

Abstract Cognitive Maps for Complex Social Systems

Oriel FeldmanHall[1,2], Jae-Young Son[1], and Apoorva Bhandari[1]

[1]Department of Cognitive and Psychological Sciences, Brown University, and
[2]Carney Institute for Brain Science, Brown University

Abstract

How do people represent social networks in their minds? Inspired by work on spatial navigation, recent research reveals how people use domain-general computational principles to build cognitive maps for navigating their social environments. However, some aspects of our social worlds, such as the densely interconnected networks we are embedded in—and the dynamics of information flow within them—challenge the particular construct of a Euclidean cognitive map that has evolved in the study of spatial navigation. Recent research reveals different types of abstract representations people can use to build efficient cognitive maps for navigating social networks. We argue that to solve challenges inherent to navigating social relationships (e.g., figuring out whom to trust or gossip with, building coalitions made up of weak ties), people build cognitive maps of both the direct and indirect relational ties surrounding them. Although the incorporation of indirect ties makes these maps nonveridical, their addition aids in flexible, adaptive behavior, which can be used for successfully navigating any complex social environment.

Keywords

cognitive maps, social networks, abstraction, information flow, link prediction, replay

Two friends, Henri and Isabel, are competing for the same prestigious award. Henri's biting remarks about Isabel in the presence of colleagues ignite hushed, backroom conversations about his lack of collegiality. Before long, gossip spreads, putting Henri's reputation at risk. As this example highlights, the joys, conflicts, and intrigues of human life unfold within our social networks, in which countless individual fates are deeply intertwined. These social networks, with their myriad connections, serve as the original information superhighways, efficiently transmitting norms, ideas, news, and influence. Every social action—whether sharing gossip or brokering a friendship—can trigger far-reaching consequences that can unexpectedly ripple across the network. The ability to navigate these social networks therefore lies at the core of social cognition.

We define social navigation as the process of making strategic choices to enhance one's relative position in a social network and achieve social goals. These strategic behaviors come in many flavors: A job applicant might rely on "weak ties" to seek out those who are well-connected to hiring managers. An affable social butterfly might organize dinner parties that foster community within a group of newcomers. Someone with juicy gossip might strategically withhold it from a loose-lipped individual to keep the gossip contained. These strategic behaviors, whatever their proximate goal may be, can change the relative social position of individuals to one other, thus subtly altering the structure of the network. Social navigation skills are ubiquitous in everyday life and are critical for accessing and maintaining social support, which plays a positive role in fostering well-being and buffering against mental health issues (Wang et al., 2018). Despite their fundamental importance, the mental representations and cognitive mechanisms enabling adaptive social navigation in humans have not been thoroughly studied.

Social navigation poses a formidable cognitive challenge given the huge number of relationships that could exist in a network. The cognitive computations needed to effectively navigate a two-person social exchange

Corresponding Author:
Oriel FeldmanHall, Department of Cognitive and Psychological Sciences, Brown University
Email: oriel.feldmanhall@brown.edu

blossom exponentially if we consider that humans are embedded in an interconnected social world (FeldmanHall & Nassar, 2021). Despite the difficulty of making informed judgments about vast social networks, it has long been known that people are remarkably adept at doing so (Travers & Milgram, 1969). To understand the cognitive basis of social navigation, we focus on two common computational problems that humans routinely solve in the social world. The first is the problem of *social link prediction*, which involves relying on patchy observations of social interactions in our network to infer the existence of unobserved relationships or to predict new relationships that may form in the future (Son et al., 2023). The second is the problem of predicting *how information flows* between people in a social network. Understanding how people tackle these two fundamental problems is crucial for understanding social navigation writ large.

Whether predicting social links or the flow of information, people must rely on an internal map of the social network, built from the social interactions they observe or hear about. Neuroscience research on spatial navigation offers a theoretical window into how we learn and organize complex, relational information. It is well established that knowledge about the underlying structure of the physical environment is encoded in a mental representation known as a "cognitive map," which specifies how locations in a physical space are related to each other (O'Keefe & Nadel, 1978; Tolman, 1948). Cognitive maps afford efficient inference about the relational structure of environments, enabling adaptive and flexible navigation (Tolman, 1948).

Inspired by research on spatial navigation, recent work has begun investigating social cognitive maps (Park et al., 2020; Peer, Hayman, et al., 2021; Tavares et al., 2015). What do they look like? What are they good for? How are they built? We review emerging research describing how the systems involved in representing the physical environment are also involved in encoding various types of cognitive maps of the social environment. We then examine how these maps can be used to solve a host of social navigation problems, from reasoning about possible relationships in the network (Son et al., 2021, 2023), tracking the flow of gossip (Xia et al., 2024), and inferring social position (Tavares et al., 2015) to figuring out how to climb the social ladder into positions of great influence (Aslarus et al., in press). We also highlight how social cognitive maps differ fundamentally from spatial maps.

The Psychological Landscape of Social Cognitive Maps

Historically, cognitive maps have mainly been studied in the domain of spatial navigation, in which evidence of such maps has been found in the medial temporal lobe of the brain (O'Keefe & Nadel, 1978). Given that spatial navigation involves tracking your position on a two-dimensional (2D) surface, the dominant idea in the literature from this long line of work is that of a cognitive map encoding a 2D Euclidean space. With a 2D map of a physical environment, you can figure out how to get from Point A to Point B—the Arc de Triomphe to a small gallery in the Marais—even if you have not made that particular journey before.

Recent work has generalized this idea of a 2D Euclidean cognitive map to social domains (Park et al., 2020; Tavares et al., 2015), showing that social maps encode a person's position in an abstract space defined by continuous social variables such as power and competence. This work has found that the same neural machinery in the medial temporal lobe that encodes spatial maps also constructs abstract maps of social variables. Although such a map is undoubtedly useful, the structure of our social environment is much too complex to be captured by metric variables. People's social "position" is determined by their place in a non-Euclidean social network, in which myriad relationships (of all sorts: family, friend, colleague) shape their access and influence. Simply put, Euclidean coordinates cannot define relationships between people in a social network, rendering it impossible to travel "east of Edith."

Our social environments are better described as a set of links (e.g., friendships; Fig. 1a) between network members, a format sometimes referred to as a graph (Peer, Brunec, et al., 2021; Peer, Hayman, et al., 2021). Such a graph encodes nonlinear spaces; comprises structures such as communities, hubs, and brokers; and is a natural analogue to people's reports of network structure in anthropological and sociological studies (Brands, 2013; Krackhardt, 1987). In highly stereotyped networks such as family trees, graph-like representations enable rapid inference of specific relations (e.g., my mother's sister is my aunt) that are not learned explicitly (Whittington et al., 2020). Therefore, a graph-like format provides a suitable candidate for the format of social cognitive maps.

Formats of Cognitive Maps for Social Spaces

One possibility is that social cognitive maps are simply a veridical representation of the many relationships in a network. Mathematically, this can be expressed in the form of an *adjacency matrix*, which represents the true relationships between all individuals (in which each cell in the matrix is a number indicating the existence of friendship—or not—between two individuals). However, representing reality as faithfully as possible may be feasible only when cognitive resources are limitless

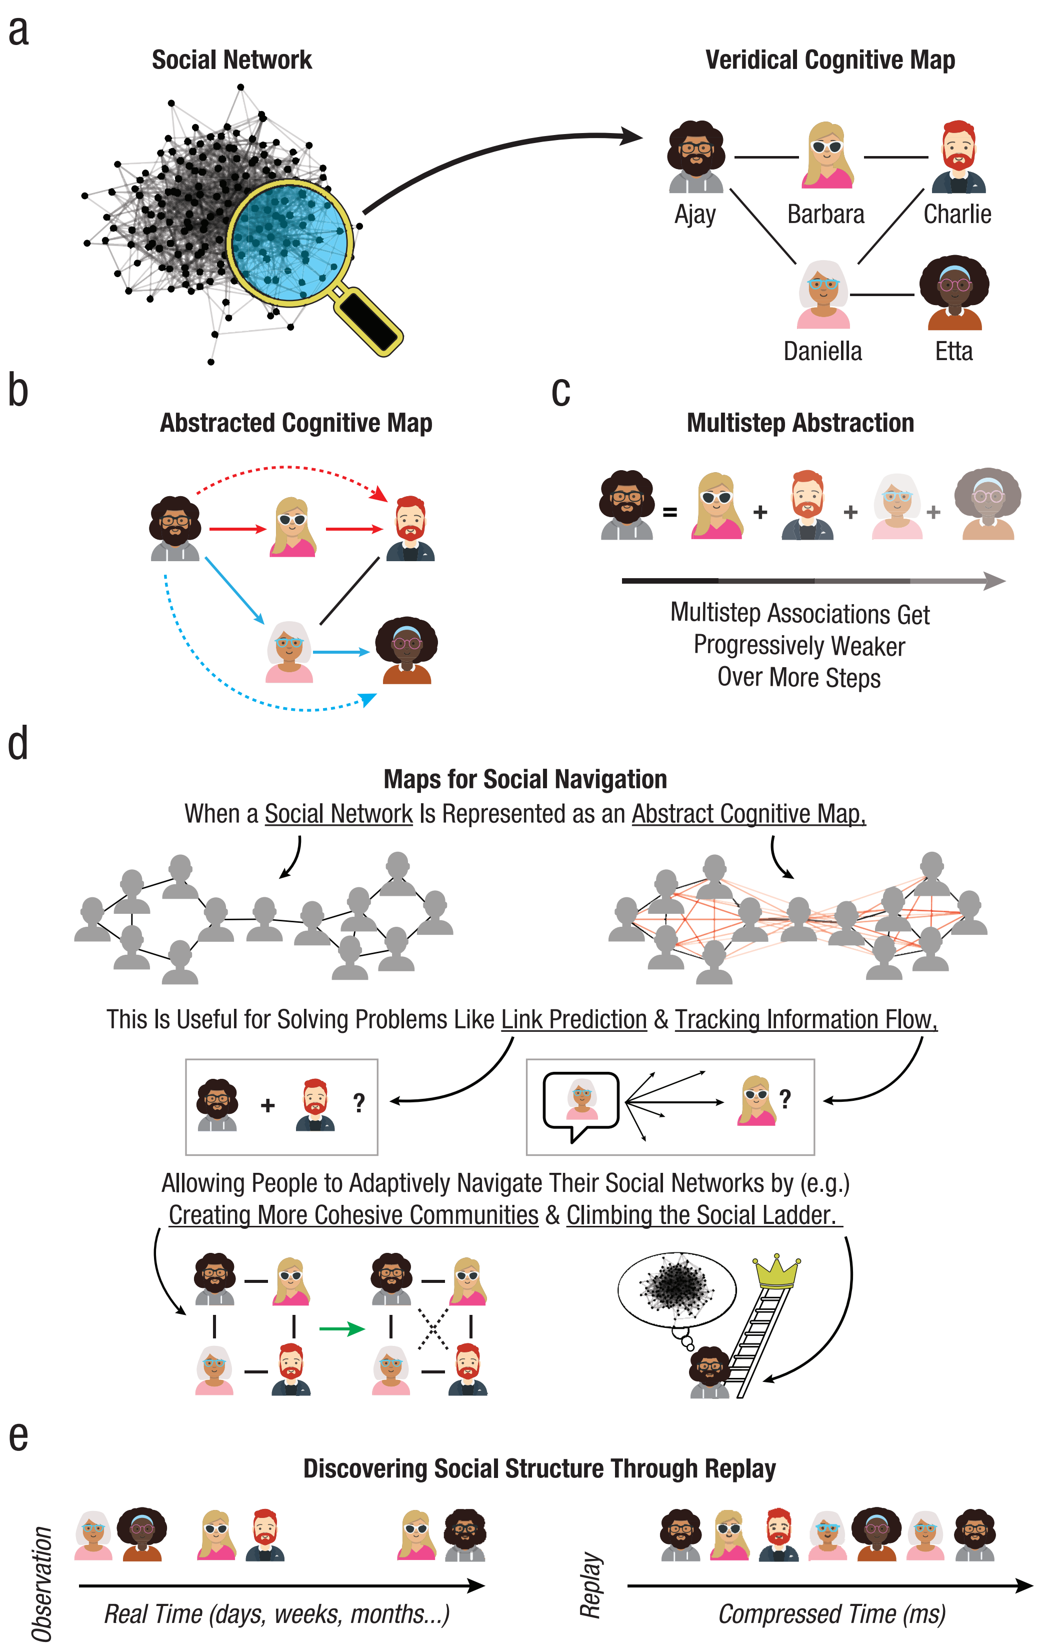

Fig. 1. *(continued on next page)*

Fig. 1. Abstract social cognitive maps. One way to represent a social network is to (a) build a veridical cognitive map consisting only of direct connections between people (e.g., knowing that Ajay and Daniella are friends). However, the large number of relationships in most real-world networks render this strategy infeasible. An alternative is to (b) represent abstracted cognitive maps consisting of both direct connections (friendships) and indirect connections (e.g., friends-of-friends and longer-range connections). This allows individuals to make informed guesses about the existence of social relationships that they may not have observed directly (e.g., inferring the likelihood that Ajay and Etta are friends despite never having seen them together). Recent research has begun to characterize what kinds of cognitive and computational processes are used to represent abstracted cognitive maps. One particularly important mechanism is (c) multistep abstraction, in which an individual is represented as the weighted sum of their friends, weighted sum of their friends, friends-of-friends, friends-of-friends-of-friends, and so on. By integrating over multiple steps, the resulting representation enables probabilistic inference of social relationships. Maps for (d) social navigation are useful for solving problems, such as predicting links (i.e., inferring the existence of unobserved relationships) and tracking information flow (e.g., predicting how gossip will spread through a network). This, in turn, allows people to better navigate their social networks, including creating more cohesive communities and climbing to more influential roles in the network. It is challenging to build an abstract cognitive map of multistep relations from direct experience alone because people must learn from piecemeal, disjointed observations of others' social relationships over long periods of time. One way the brain can overcome this problem is by (e) "replaying" temporally compressed sequences of social relationships. As representations of Ajay and Charlie are now brought closer together in time, it is easier to learn their multistep links. All avatar icons were generated from https://getavataaars.com.

and memory is perfect. As the size of a network increases, the number of possible relationships explodes combinatorically, severely taxing our ability to learn, remember, and use such a representation to do anything useful, such as tracking information flow (Breza et al., 2018). Moreover, social networks are constantly evolving as relationships are made and broken. Therefore, memorization may not be a pragmatic option in most dynamic and complex social settings.

For decades, social psychologists have recognized that people do not encode veridical information about social ties. Instead, their memories are prone to error and systematically biased, inventing new relationships where none exist and forgetting about relationships that do exist. In other words, social network representation is not precise. Past accounts of this error-prone phenomenon have posited that people use heuristics or "schemas" because of the need to store knowledge in a compressed form given memory limits (Brashears, 2013) or to reveal enduring patterns in relationships (Brands, 2013; Krackhardt, 1987). According to these theories, people's representations are fuzzy because they simply cannot encode accurate knowledge about such a large number of possible connections and instead focus on just the important features. We offer an alternate perspective: Social cognitive maps are fuzzy—or abstract—so that they can be predictive.

More than 70 years ago, Festinger (1949) realized that if you start with an adjacency matrix, simple matrix multiplication could be used to compute information about an individual's long-range connections in the network—including indirect ones, such as two individuals connected only through a chain of intermediary friends. This deceptively simple mathematical insight can be used to create an abstract psychological map reflecting how much "influence" network members have over direct and indirect ties (Katz, 1953). Such a map is well suited to solve our core problems of predicting information flow and unobserved social links.

To illustrate, imagine that an individual shares a rumor with friends. The rumor is most likely to spread to friends of friends and a little less likely to spread to friends of friends of friends, and so on. Mapping long-range indirect ties makes it easier to predict the flow of information. In a recent experiment, people were tasked with spreading gossip widely without the target of the gossip finding out. We used computational modeling to identify the mental representation of the network used to make gossip decisions (Xia et al., 2024). Rather than encoding each individual friendship (a veridical adjacency matrix), we discovered that people represent the strength of connectivity between individuals on the basis of both direct (Fig. 1b, solid lines) and indirect (Fig. 1b, dotted lines) ties, with indirect ties being discounted by the length of the indirect path. This is termed "multistep abstraction" (Son et al., 2023; Fig. 1c), which enables the likelihood that information will pass between individuals in the network to be quantified.

Cognitive Maps for Social Navigation Problems

Abstract social maps capture information about the global structure of a network, enabling inferences about how information flows through multiple channels simultaneously (Xia et al., 2024; Fig. 1d). By encoding complex social environments in this format, we gain predictive power at the cost of accuracy, which helps explain how humans can achieve the spectacular feat of near-constant gossiping while ensuring that sensitive information does not fall into the wrong hands (Blumberg, 1972). In a separate study that probed how the brain achieves this feat in large real-world social networks, we found that the medial temporal lobe spontaneously encodes cognitive maps containing information about connectivity between individuals, which is shaped by both the direct and indirect ties of the network (Teoh et al., 2025). Moreover, greater access to these maps predicts the subsequent successful brokerage of connections between immediate friends, which makes one's local clique more socially cohesive (Fig. 1d).

An abstract map of indirect ties also helps identify individuals in the broader network that one is more likely to encounter, either through friends or friends-of-friends. Therefore, maps also help solve social link-prediction problems (Fig. 1d). In a series of experiments with artificial social networks learned in the lab, we recently demonstrated that rather than relying on an adjacency matrix, people solve link prediction problems by building an abstract cognitive map that integrates both direct and indirect connections (Son et al., 2023). This enables the discovery of probabilistic links between network members who have never been observed interacting—but are likely to exist given the structure of the network. Indeed, in ongoing work analyzing data from a longitudinal study of an evolving, real-world social network, we have found evidence that people can make accurate predictions about who will become friends (and which friendships will dissolve) up to 6 months in advance of the tie formation (or dissolution). Our research shows that fuzziness is an asset—not a liability—when navigating and reasoning about social relationships, contrary to what past theories have argued.

There are other kinds of abstract cognitive maps that people can construct. Maps need not represent the actual relationships that make up the network. In fact, by abstracting over individuals altogether, people can use socially relevant features, such as group membership, profession, hobbies, and so forth, to make inferences about the likelihood of existing social ties (Son et al., 2021). This affords the ability to solve a very different kind of link prediction problem. If you can learn the statistics of the social world, even when social ties are not based on obvious homophily (e.g., vegans are often friends with marathon runners—two groups who enjoy detailing their accomplishments out loud), then you can generalize this mapping to any scenario or social environment with ease. This logic extends to any patterns that exist in a social network—such as the likelihood that a friend's enemy is an enemy. By tracking the natural statistics of relationship patterns, people can make inferences about positive and negative ties in a signed social network (Hu et al., 2025).

Together, this recent body of work demonstrates that people can encode a variety of abstracted cognitive maps of complex social spaces. We have argued that the format most suitable for social navigation is that of a multistep, graph-like cognitive map that integrates both direct and indirect ties. Indeed, such a map is spontaneously encoded in the medial temporal lobe (Teoh et al., 2025) and contributes to various forms of social navigation (Aslarus et al., in press; Hu et al., 2025; Son et al., 2021, 2023, 2024; Xia et al., 2024; Teoh et al., 2025). This notion of a cognitive map is quite different from the dominant idea of a Euclidean cognitive map because it is shaped both by the unique features of social networks and the demands of social navigation. For example, information flow has very different dynamics compared with moving through space. As information (e.g., gossip) travels, it is copied and transmitted simultaneously along many paths. The cognitive maps with multistep connectivity that we observe capture exactly this type of dynamic. These unique features make social maps distinct from spatial maps and challenge the dominant construct of a 2D Euclidean cognitive map.

Discovering Structure in Social Networks During Sleep

How does an individual use observations of their social world, such as seeing two people having lunch together, to build a cognitive map containing multistep ties? To build a map, people must repeatedly explore their environment to understand its broader relational structure. Spatially, this is akin to moving into a new neighborhood and taking a walk to get the "lay of the land." Although observations in our social worlds provide information only about direct ties, people can use a mechanism such as association chaining (Kumaran & McClelland, 2012; Lynn et al., 2020; Momennejad et al., 2017; Son et al., 2023, 2024) to infer longer-range connections from observed social interactions. However, this poses a significant cognitive challenge because building up a full map of a network requires a very large number of observations.

Moreover, although humans can explore a physical environment by moving through space, it is information that travels through social networks. Information flow trajectories are often not observable and must be inferred from partial cues, such as information revealed through a conversation. In other words, to get the lay of the land for a social network, we must explore it in our minds, simulating how information might flow through ties that we know about and others that we might infer. How might we do this efficiently? We turn again to spatial navigation research for inspiration. During rest periods, such as when a rat is taking a break from exploring a maze, the brain spontaneously "replays" the experience of running around the environment. Replay, in its simplest form, is characterized by sequential re-instantiation of actual experiences in a temporally compressed fashion, both during awake rest and slow-wave sleep (Foster, 2017). Replay of information trajectories through social networks, combined with associative chaining, allows individuals to stitch together a picture of the network's global structure.

However, given that most of our observations comprise interactions between individuals with direct ties, and longer-range information flow trajectories are very rarely observed, what is the source of this replay? A possible solution is for the brain to instantiate activity patterns reflecting sequences that have never been experienced before but are consistent with known ties in the network, a phenomenon known as *generative replay* (Stoianov et al., 2022; Fig. 1e). In other words, to build a more comprehensive cognitive map, the brain—while at rest—supplements direct experience by synthetically generating entirely new experiences that are consistent with the known structure of the network (Liu et al., 2021; Momennejad et al., 2018). In a social network, this may be instrumental in building a useful map of the network. Indeed, we have found that a replay-like process during overnight rest helps people track how information might flow through a social network, resulting in inferences that can span communities (Son et al., 2024).

This research suggests that extended periods of rest, such as sleep, might be instrumental for building the kinds of abstract representations needed for solving information-tracking or link prediction problems. Research in other domains illustrates that replay during sleep supports the consolidation of memory traces into a broader framework of existing memories (Feld & Born, 2017), which helps transform initial experiences into an abstracted representation that can be used flexibly across contexts (Farzanfar et al., 2023; Lewis & Durrant, 2011). Sleep has consistently been associated with the process of abstraction (Lewis & Durrant, 2011; Walker & Stickgold, 2010), hinting at the role of sleep-associated replay in building abstract maps. Abstraction critically reveals the underlying structure of a given environment, essentially carving the environment at its joints. Applied to a social network, this is akin to bringing into relief the hubs, cliques, and brokers (bridges between otherwise disconnected cliques) that serve as the major pipelines through which information can flow. To prioritize these core structural elements of the network, it is likely that replay during sleep facilitates the abstraction of relational information into a more generalized map.

Conclusions

Being able to predict the consequences of our actions while embedded in a web of connections is fundamental for adaptively navigating our social world. Borrowing ideas from spatial navigation research, recent work shows how people use domain-general computational principles to build cognitive maps of social networks. At the same time, the unique features of social networks challenge traditional assumptions that originated from a historical focus on spatial navigation. We endorse the emerging consensus that cognitive maps are useful across many domains (Behrens et al., 2018) and argue that building abstract social maps can help make the large space of possible relationships cognitively tractable (Tenenbaum et al., 2011). Flexible, adaptive behavior is then achieved by drawing on these abstract cognitive maps, often in ways that systematically deviate from an individual's direct observations of reality. Indeed, we have found that building abstract mental representations that prioritize structural features of one's own social network, such as cliques or groups (at the expense of knowing about specific relationships), predicts whether an individual will climb the social ladder to become influential (Aslarus et al., in press). In contrast, those who fail to abstract over relationships to detect underlying structure, even if once influential, tend to lose their social prominence over time. This suggests that these fuzzy representations are functionally adaptive for inferring relationships (Son et al., 2021, 2023, 2024; Xia et al., 2024), a perspective we believe will buy enormous predictive power in understanding how humans navigate their social worlds.

Recommended Reading

Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). (See References). Broadly reviews the computational ideas that inform current thinking about cognitive maps, including ideas from research in spatial navigation, reinforcement learning, and neuroscience.

Son, J.-Y., Bhandari, A., & FeldmanHall, O. (2023). (See References). Demonstrates that errors in people's memories of social relationships in a network result from learning a cognitive map that integrates direct and indirect ties with multistep abstraction.

Son, J.-Y., Vives, M.-L., Bhandari, A., & FeldmanHall, O. (2024). (See References). Demonstrates that people's cognitive map of a social network becomes more abstract

overnight and that their ability to engage in long-range social navigation is enhanced by overnight rest.

Xia, A., Teoh, Y. Y., Nassar, M., Bhandari, A., & FeldmanHall, O. (2025). (See References). Demonstrates that people strategically share gossip in social networks by leveraging internal models prioritizing popularity and distance to assess how information might spread.

Transparency

Action Editor: Robert L. Goldstone
Editor: Robert L. Goldstone
Declaration of Conflicting Interests
The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.
Funding
This work is supported by the National Science Foundation award 2123469 (O.F.H. and A.B.).

ORCID iDs

Oriel FeldmanHall https://orcid.org/0000-0002-0726-3861
Jae-Young Son https://orcid.org/0000-0002-3852-1237
Apoorva Bhandari https://orcid.org/0000-0001-5295-1623

Acknowledgments

We would like to thank an anonymous reviewer whose feedback helped us develop a useful, general definition of social navigation.

References

Aslarus, I., Son, J. Y., Xia, A., & FeldmanHall, O. (in press). Early insight into social network structure predicts climbing the social ladder. *Science Advances*.

Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*, *100*(2), 490–509.

Blumberg, H. H. (1972). Communication of interpersonal evaluations. *Journal of Personality and Social Psychology*, *23*(2), 157–162.

Brands, R. A. (2013). Cognitive social structures in social network research: A review. *Journal of Organizational Behavior*, *34*(Suppl. 1), S82–S103.

Brashears, M. E. (2013). Humans use compression heuristics to improve the recall of social networks. *Scientific Reports*, *3*, Article 1513. https://doi.org/10.1038/srep01513

Breza, E., Chandrasekhar, A. G., & Tahbaz-Salehi, A. (2018). *Seeing the forest for the trees? An investigation of network knowledge* (Working Paper No. 24359). National Bureau of Economic Research. https://www.nber.org/papers/w24359

Farzanfar, D., Spiers, H. J., Moscovitch, M., & Rosenbaum, R. S. (2023). From cognitive maps to spatial schemas. *Nature Reviews Neuroscience*, *24*(2), 63–79.

Feld, G. B., & Born, J. (2017). Sculpting memory during sleep: Concurrent consolidation and forgetting. *Current Opinion in Neurobiology*, *44*, 20–27.

FeldmanHall, O., & Nassar, M. R. (2021). The computational challenge of social learning. *Trends in Cognitive Sciences*, *25*(12), 1045–1057.

Festinger, L. (1949). The analysis of sociograms using matrix algebra. *Human Relations*, *2*, 153–158.

Foster, D. J. (2017). Replay comes of age. *Annual Review of Neuroscience*, *40*(1), 581–602.

Hu, Y.-F., Vives, M. L., Bhandari, A., & FeldmanHall, O. (2025). *Using statistical learning to navigate signed social network*. PsyArXiv. https://doi.org/10.31219/osf.io/2fpt5_v1

Katz, L. (1953). A new status index derived from sociometric analysis. *Psychometrika*, *18*(1), 39–43.

Krackhardt, D. (1987). Cognitive social structures. *Social Networks*, *9*(2), 109–134.

Kumaran, D., & McClelland, J. L. (2012). Generalization through the recurrent interaction of episodic memories: A model of the hippocampal system. *Psychological Review*, *119*(3), 573–616.

Lewis, P. A., & Durrant, S. J. (2011). Overlapping memory replay during sleep builds cognitive schemata. *Trends in Cognitive Sciences*, *15*(8), 343–351.

Liu, Y., Mattar, M. G., Behrens, T. E. J., Daw, N. D., & Dolan, R. J. (2021). Experience replay is associated with efficient nonlocal learning. *Science*, *372*(6544), Article eabf1357. https://doi.org/10.1126/science.abf1357

Lynn, C. W., Kahn, A. E., Nyema, N., & Bassett, D. S. (2020). Abstract representations of events arise from mental errors in learning and memory. *Nature Communications*, *11*(1), Article 2313. https://doi.org/10.1038/s41467-020-15146-7

Momennejad, I., Otto, A. R., Daw, N. D., & Norman, K. A. (2018). Offline replay supports planning in human reinforcement learning. *eLife*, *7*, Article e32548. https://doi.org/10.7554/eLife.32548

Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N. D., & Gershman, S. J. (2017). The successor representation in human reinforcement learning. *Nature Human Behaviour*, *1*(9), 680–692.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Clarendon Press.

Park, S. A., Miller, D. S., Nili, H., Ranganath, C., & Boorman, E. D. (2020a). Map making: Constructing, combining, and inferring on abstract cognitive maps. *Neuron*, *107*(6), 1226–1238.e8. https://doi.org/10.1016/j.neuron.2020.06.030

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.

Peer, M., Hayman, M., Tamir, B., & Arzy, S. (2021). Brain coding of social network structure. *Journal of Neuroscience*, *41*(22), 4897–4909.

Son, J.-Y., Bhandari, A., & FeldmanHall, O. (2021). Cognitive maps of social features enable flexible inference in social networks. *Proceedings of the National Academy of Sciences, USA*, *118*(39), Article e2021699118. https://doi.org/10.1073/pnas.2021699118

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, USA*, *120*(47), Article e2310801120. https://doi.org/10.1073/pnas.2310801120

Son, J.-Y., Vives, M.-L., Bhandari, A., & FeldmanHall, O. (2024). Replay shapes abstract cognitive maps for efficient social navigation. *Nature Human Behaviour*, *8*, 2156–2167. https://doi.org/10.1038/s41562-024-01990-w

Stoianov, I., Maisto, D., & Pezzulo, G. (2022). The hippocampal formation as a hierarchical generative model supporting generative replay and continual learning. *Progress in Neurobiology*, *217*, Article 102329. https://doi.org/10.1016/j.pneurobio.2022.102329

Tavares, R. M., Mendelsohn, A., Grossman, Y., Williams, C. H., Shapiro, M., Trope, Y., & Schiller, D. (2015). A map for social navigation in the human brain. *Neuron*, *87*(1), 231–243.

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, *331*(6022), 1279–1285. https://doi.org/10.1126/science.1192788

Teoh, Y. Y., Son, J.-Y., Xia, A., Bhandari, A., & FeldmanHall, O. (2025). *Medial temporal lobe encodes cognitive maps of real-world social networks*. bioRxiv. https://doi.org/10.1101/2025.03.18.643925

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*(4), 189–208.

Travers, J., & Milgram, S. (1969). An experimental study of the small world problem. *Sociometry*, *32*(4), 425–443.

Walker, M. P., & Stickgold, R. (2010). Overnight alchemy: Sleep-dependent memory evolution. *Nature Reviews Neuroscience*, *11*(3), 218–218.

Wang, J., Mann, F., Lloyd-Evans, B., Ma, R., & Johnson, S. (2018). Associations between loneliness and perceived social support and outcomes of mental health problems: A systematic review. *BMC Psychiatry*, *18*, Article 156. https://doi.org/10.1186/s12888-018-1736-5

Whittington, J. C. R., Muller, T. H., Mark, S., Chen, G., Barry, C., Burgess, N., & Behrens, T. E. J. (2020). The Tolman-Eichenbaum machine: Unifying space and relational memory through generalization in the hippocampal formation. *Cell*, *183*(5), 1249–1263.e23.

Xia, A., Teoh, Y. Y., Nassar, M., Bhandari, A., & FeldmanHall, O. (2024). *Knowledge of information cascades through social networks facilitates strategic gossip*. PsyArXiv. https://doi.org/10.31234/osf.io/yq82d