Replay shapes abstract cognitive maps for efficient social navigation

Jae-Young Son[1*], Marc-Lluís Vives[2*], Apoorva Bhandari[1**], & Oriel FeldmanHall[1,3**]
[1]Department of Cognitive, Linguistic, and Psychological Sciences, Brown University, Providence, RI, USA
[2]Institute of Psychology, Leiden University, Leiden, The Netherlands
[3]Robert J. and Nancy D. Carney Institute for Brain Science, Brown University, Providence, RI, USA

* Equal contribution
** Equal contribution

Corresponding Authors:
Apoorva Bhandari
Brown University
190 Thayer St. Providence, RI 02912
apoorva_bhandari@brown.edu

Oriel FeldmanHall
Brown University
190 Thayer St. Providence, RI 02912
oriel.feldmanhall@brown.edu

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.

MAIN

Human social life is embedded within a complex web of connections where the consequence of a social misstep can reverberate far beyond the original interaction. For example, trusting the wrong person with a secret may result in embarrassing information spreading past one's immediate circles to distant individuals and communities. Navigating through the social world therefore requires representing people's relationships with one another, including those that fall outside of one's direct circle of friends[1-3]. Little is known, however, about the mental representations that enable social navigation through complex networks, nor how these representations are built.

What kinds of mental representations might support social navigation? Decades of research on spatial navigation offers a useful window into how humans might organize complex relational information. It is well-established that knowledge about physical environments is represented in cognitive maps of spatial relationships[4-6]. The format of these spatial maps allows objects to be placed within two-dimensional mental spaces[5,7], affording representation of the longer-range spatial relationships between those objects. Outside of spatial navigation, recent work demonstrates that humans also represent abstract maps of conceptual spaces[8-10], including social traits such as competence and popularity[11,12]. However, the relationships in social networks are poorly characterized by two-dimensional spaces, and it is not known what alternative format(s) might instead be used to build abstract cognitive maps of social networks.

Recent work in cognitive neuroscience points to a promising representational format for social networks. Network representations can be built through multistep abstraction—a mechanism borrowed from the Successor Representation framework in reinforcement learning—which encodes not only the direct connections between entities (e.g., friendships), but also longer-range, multistep connections (e.g., friends-of-friends)[10,13-19]. By adjusting how many steps are integrated over, network representations can be learned at various levels of abstraction[20], where greater abstraction confers rapid inference about distant relations, as well as network structures like communities[3,14,21]. This ability to represent longer-range relations likely aids social navigation, including tasks such as predicting where gossip might spread if shared with a given individual.

A second question revolves around how people efficiently build maps from limited direct experience. Evidence from rodent neuroscience, and more recently humans, points to an important role of replay, where the brain generalizes from experience to simulate synthetic sequences, especially prioritizing those that are most critical for adaptive navigation[22-28]. Indeed, it has long been noted that a replay-like mechanism appears necessary to learn sufficiently abstract representations for navigation[14-16]. Offline replay during sleep appears to play an especially important role in generating more abstract representations of the environment[29-33]. As abstraction can help reveal the underlying structure of a given environment and therefore aid longer-range navigation[14,17], it is likely that sleep is more critical for building the kinds of abstract representations needed for longer-range navigation through social networks.

Multistep abstraction therefore not only specifies a useful format for representing the topology of graph structures like social networks[3,10,15-17,21], but it also provides a natural interface between abstract cognitive maps and offline replay. Although multistep abstraction is an attractive model of how people represent and navigate social networks, past research has only established that people's memory representations of social networks are consistent with multistep abstraction, and it is yet unknown whether or how multistep abstraction supports navigation behaviors.

Here 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. To assess whether humans use cognitive maps to solve the challenge of social navigation, we created a task where subjects learn about friendships in a social network (Figure 1A), allowing us to probe whether they could navigate how information flows through the community (Figure 1C). We then either had subjects take the navigation task immediately, or brought subjects back to the laboratory the next day after overnight rest to test whether navigation accuracy improved after sleep (Figure 1F). Using computational modeling, we characterized the underlying cognitive maps employed by subjects and further tested whether a replay-like mechanism helps to scaffold more successful social navigation.

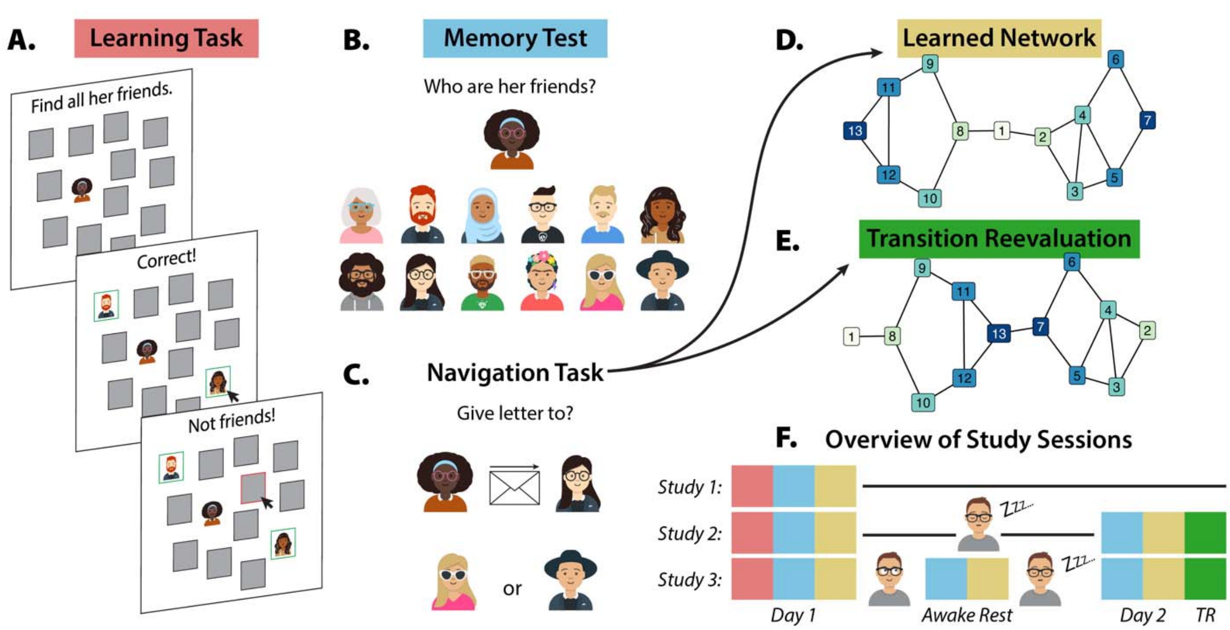

Figure 1 | Study design. A. *The learning task used 'flash cards' to facilitate rapid and accurate learning. When presented with a Target network member, participants were required to find all of the Target's friends from the face-down cards. When a card corresponding to the Target's friend was selected, the card flipped face-up to reveal their photograph. Cards remained face-down in response to incorrect guesses.* ***B.*** *The memory test presented a Target and required participants to indicate all of the Target's friends. No feedback about accuracy was ever provided.* ***C.*** *The social navigation task presented a network member wishing to send a message to a Target through one of two Sources. Participants were required to indicate which Source was the better choice for efficient delivery to the Target.* ***D.*** *The social network learned by participants.* ***E.*** *In Studies 2-3, participants were informed that some of the friendships had been broken, and that others had formed. This necessitated rapid reevaluation of how network members were related to each other.* ***F.*** *A schematic illustrating what tasks participants completed on what days, in which studies. The color-coding corresponds to the task labels in parts A-E of this figure; yellow and green indicate completion of the navigation task for the learned and reevaluated networks, respectively. All avatar icons were generated using getavataaars.com, designed by Pablo Stanley and developed by Fang-Pen Lin.*

Humans can efficiently solve social navigation problems.

We developed a novel 'message-passing' task as an experimental testbed of flexible social navigation, which assessed whether people understand how information flows throughout the network (Figure 1C). On each trial, a network member wished to pass a letter to a Target within the network, and needed to choose between Sources A and B. 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 was delivered to the Target. The subject's task was to choose the Source that would result in the most efficient delivery. Trials were classified according to the shortest path distance from the network member who had sent the letter. For example, when the correct Source was directly friends with the Target, we classified these as distance-2 problems, as the message needed to be passed twice to reach the Target. An accurate response was defined as choosing the Source with the shortest path to the Target (Methods). The Target changed from trial-to-trial, such that successful navigation required flexible use of knowledge about connections between network members. To rule out the possibility that navigation accuracy might improve simply from experience with the task, subjects were never provided with feedback.

Across three laboratory studies (total N = 146; data pooled for efficiency, but results replicate across all studies), subjects completed this navigation task shortly after learning friendships in a novel social network. (Figure 1A). Subjects never observed the whole network and were provided no direct information about multistep, longer-range connections between network members (e.g., friends-of-friends), but instead only observed dyadic relationships. Despite this learning format, subjects achieved above-chance navigation accuracy not only for problems where the Source was directly friends with the Target (80% accuracy at distance-2 β = 1.68, Z = 14.30, 95% CI = [1.45, 1.91], p < .001), but also for the longer-range problems (70% accuracy at distance-3, β = 1.06, Z = 9.34, 95% CI = [0.84, 1.28], p < .001; 63% accuracy at distance-4, β = 0.66, Z = 6.05, 95% CI = [0.44, 0.87], p < .001; Figure 2A). These results suggest that subjects learned a cognitive map that supported flexible, long-range social navigation, despite only being provided pairwise information about friendships in the network.

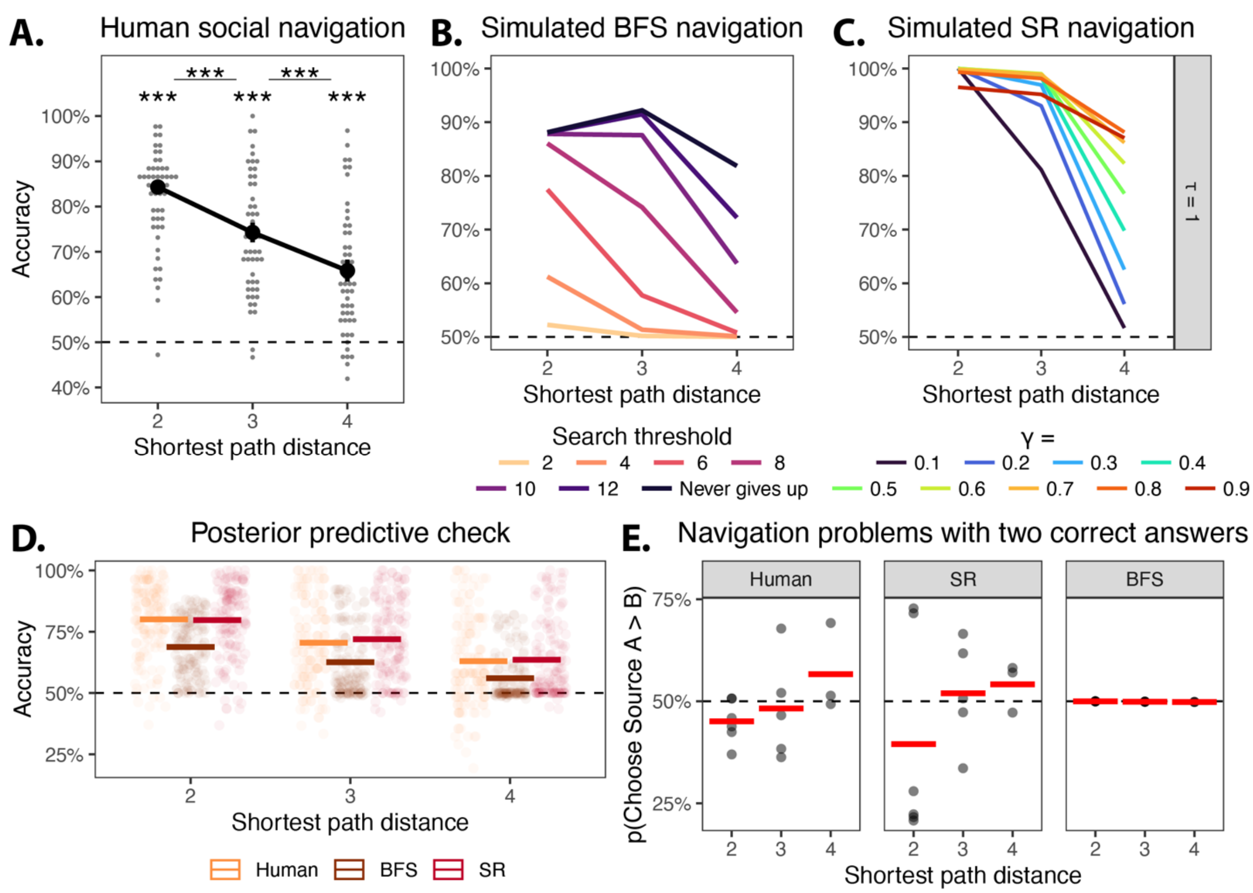

Figure 2 | Evidence of social navigation. A. *Shortly after learning about friendships in a novel social network, people are able to solve social navigation problems with above-chance accuracy. Trend lines reflect the estimated means from a mixed-effects logistic regression model, and error bars reflect estimated standard errors.* ***B.*** *Simulated navigation behavior from a Breadth-First Search (BFS) algorithm that uses an internal model for planning.* ***C.*** *Simulated navigation behavior from the Successor Representation (SR) algorithm, which learns multistep relations between network members over a horizon controlled by γ.* ***D.*** *The computational model of multistep abstraction accurately recapitulates human behavior, whereas the computational model of model-based planning does not.* ***E.*** *On a subset of navigation problems, subjects were* *presented with two Sources that were the same path distance away from the Target. Despite this, humans frequently demonstrated a preference for one Source over the other. The SR is able to mirror these preferences, while BFS is not. Individual dots reflect the average behaviors (of humans or models) for the trials in which A and B are matched. Red lines reflect the average across items.*

Computational models of social navigation.

We consider two distinct decision-making strategies that an agent could employ to flexibly solve novel social navigation problems. A normatively optimal agent would represent all pairwise friendships within the social network, then recursively iterate through those friendships until it computes the shortest path between a given Source-Target pair. In practice, online navigation of this kind is time-consuming and computationally costly[16], but can be made tractable in small networks by a planning algorithm like Breadth-First Search (BFS)[34]. In our implementation, we assume that the agent starts two parallel searches, one from each Source, and include a search threshold parameter characterizing an agent's tendency to 'give up' during long searches (see Methods). Simulation results confirm that when BFS agents are tasked with solving the same social navigation problems as humans, they are asymptotically capable of achieving high navigation accuracy (Figure 2B).

Alternatively, an agent could navigate more efficiently by caching (i.e., pre-computing) relevant knowledge. In the context of social navigation, it would be particularly useful to cache knowledge of individuals' longer-range, multistep connections (e.g., friends-of-friends). Recent work in cognitive neuroscience points to the Successor Representation (SR) as a useful format for encoding such cached knowledge[15-17], including cognitive maps of social networks[3]. The SR approximates the probability of transitioning from a Source to a Target in a given number of steps. A single parameter, the successor horizon γ, controls how many steps are integrated over, and therefore dictates how the SR integrates knowledge of shorter- vs longer-range connections. 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 integrates over longer-range connections (e.g., friends-of-friends-of-friends…). Once the agent has cached estimates of $p(\text{Target} \mid \text{Source}, \gamma)$, it can then decide between the two possible Sources using a softmax choice rule with inverse temperature τ, controlling how noisily the agent chooses.

To assess asymptotic performance, we provided SR agents with a large number of learning experiences (see Methods). Simulation results reveal that multistep abstraction is sufficient to achieve high navigation accuracy: higher values of γ were associated with greater navigation accuracy for longer-range problems, and SR agents achieved uniformly high navigation accuracy for the shorter-range problems regardless of γ (Figure 2C). These results therefore confirm that human subjects' social navigation decisions could, in principle, be supported by a cognitive map of multistep relationships.

We next asked which computational model provides a better explanation for human behavior in the social navigation task and found that the SR model provided a better group-level fit to the data than the BFS model, as well as a hybrid BFS-SR model (all PXP > 0.99; see Methods). Therefore, a formal comparison of computational models suggests that human behavior on the message-passing task is best explained by the use of a cognitive map containing cached knowledge of abstract, multistep relations. Posterior predictive checks further confirmed that the BFS model systematically mischaracterizes human subjects' navigation behaviors, while the SR model is largely successful in recapitulating human behavior (Figure 2D).

Finally, we tested whether each model was able to predict human behavior in a held-out subset of navigation problems (i.e., trials that were not used to fit model parameters). These trials were unique in that both Sources were the same shortest distance away from the Target, making them equally correct choices to a model-based agent. While the BFS model does not systematically favor one Source over the other in these trials (Figure 2E), humans demonstrate preferences for Sources that have multiple (relatively) short paths to the Target, which is mirrored by the SR model predictions (Figure 2E).

Social navigation improves with overnight rest.

To test whether a replay-like mechanism might result in improved social navigation after overnight rest (i.e., sleep), subjects in studies 2-3 (N = 96) completed a two-day procedure. The day after their first session, subjects returned to the laboratory and completed the social navigation task again. Results reveal that after overnight rest, subjects became significantly more accurate at solving problems across all distances (82% accuracy at distance-2, $\beta = 0.23$, $Z = 2.91$, 95% CI = [0.08, 0.39], $p = .004$; 75% accuracy at distance-3, $\beta = 0.26$, $Z = 2.99$, 95% CI = [0.09, 0.44], $p = .003$; 71% accuracy at distance-4, $\beta = 0.43$, $Z = 5.17$, 95% CI = [0.27, 0.59], $p < .001$; Figure 3A). This improvement was particularly pronounced for the longest-range distance-

4 problems, compared to the accuracy improvement for distance-2 problems (β = 0.20, Z = 2.36, 95% CI = [0.03, 0.36], p = .018).

To test whether a brief period of awake rest is sufficient to improve navigation accuracy, subjects in study 3 (N = 46) were allowed to rest for approximately 15 minutes at the end of the standard Day 1 procedure and before overnight rest (Figure 1F). After this brief awake rest period, they completed the same memory and navigation tasks again. This awake rest was not sufficient for improving navigation accuracy at the group level (all Ps > .1; Figure 3B), suggesting that a longer period of rest (or sleep) may be needed to produce significant improvements in navigation.

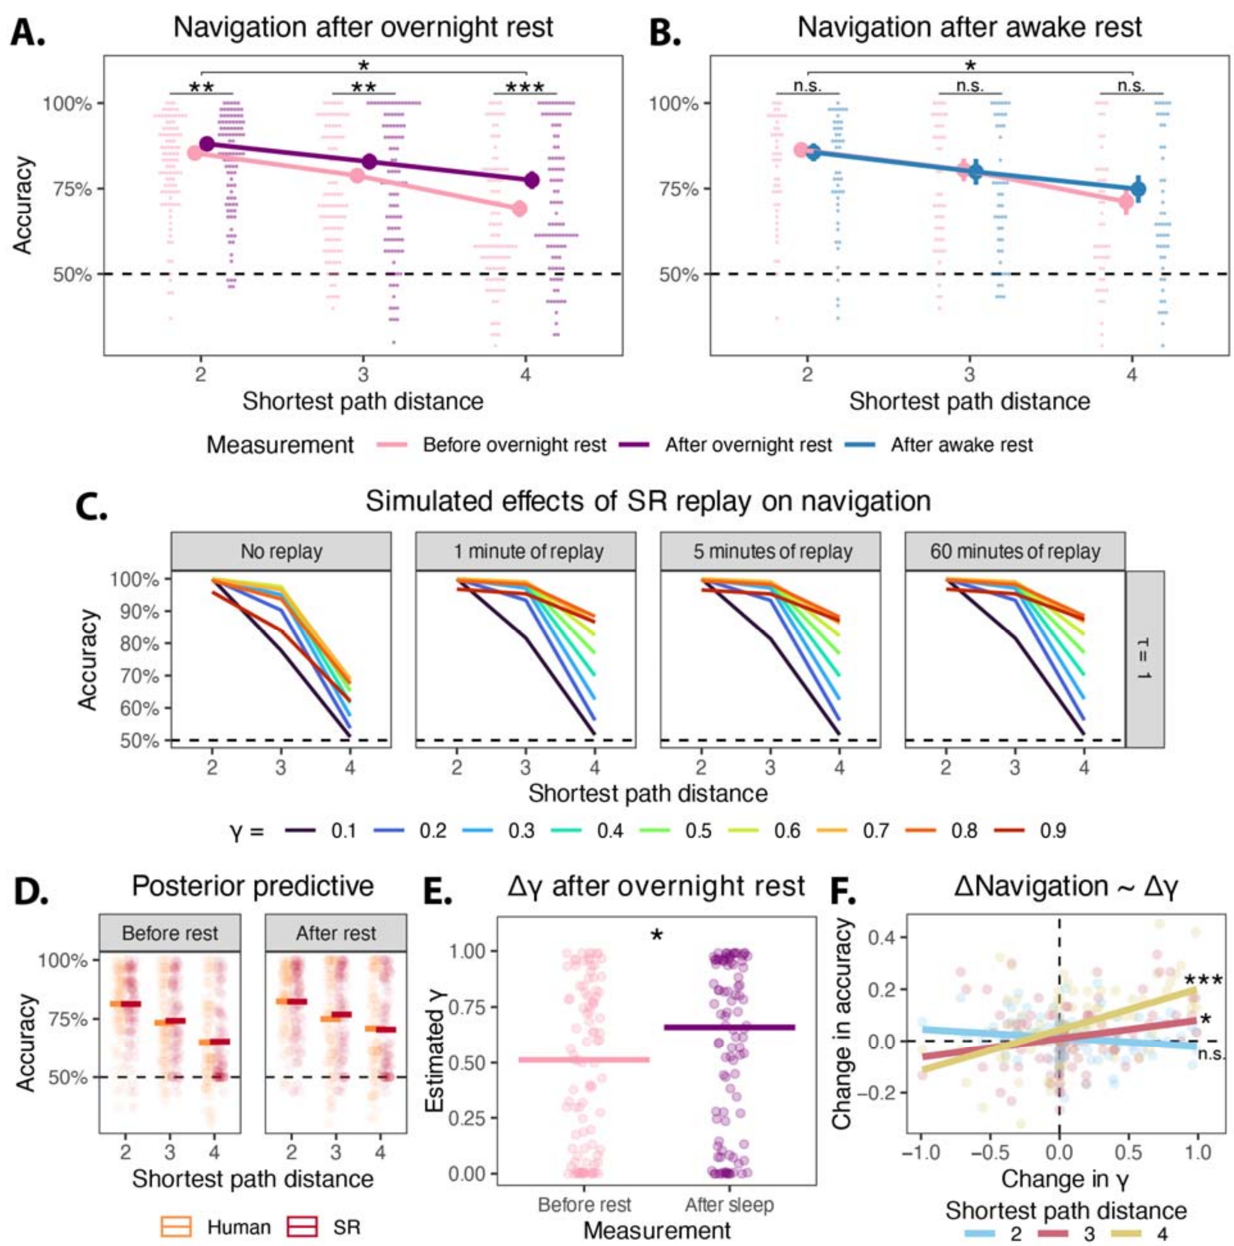

Figure 3 | Evidence for a replay-like mechanism. A. *After overnight rest, humans became more accurate at solving social navigation problems. Increases in navigation accuracy were particularly pronounced for problems requiring knowledge of longer-range connections.* ***B.*** *Brief awake rest did not significantly improve navigation. In panels A-B, trend lines reflect the estimated means from a mixed-effects logistic regression model, and error bars reflect estimated standard errors.* ***C.*** *Simulations suggest that even one minute of replay helps an agent consolidate its representation of friendships in a social network, resulting in more accurate navigation. However, consolidation rapidly asymptotes, and further replay does not result in notably improved navigation. The simulation results instead suggest that overnight rest may help an agent build more abstract representations (characterized by larger values of γ), which integrate over a greater number of multistep relations (e.g., friends-of-friends-of…), and which aid in solving longer-range navigation problems.* ***D.*** *The computational model of multistep abstraction accurately recapitulates human behavior before and after overnight rest.* ***E.***

Parameter estimates from the computational model fit to subjects' navigation behaviors. The model predicts a group-level increase in γ after overnight rest. Crossbars reflect medians. ***F.*** *Increased γ after overnight rest was associated with improved longer-range navigation, but not shorter-range navigation. Linear trend lines are shown for visualization only; the statistical tests reflect Spearman rank-correlation.*

A computational model of replay.

In the theoretical framework of the Successor Representation, replay is a natural mechanism for explaining how overnight rest improves social navigation[14-16]. 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. There are several plausible hypotheses of how replay might result in improved navigation. For example, a 'consolidation' hypothesis suggests that replay fills the gaps left by insufficient direct experience, allowing the agent to learn a more stable representation[14,23]. Alternatively, an 'abstraction' hypothesis predicts 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)[20]. 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 even more abstract cognitive maps.

To test these hypotheses, we conducted a simulation study examining how successfully an artificial agent could solve our social navigation problems, given varying amounts of replay. Drawing on past empirical studies measuring replay in both rodents and humans, our simulations assumed that it takes 50ms to replay a single transition between two states (i.e., a friendship dyad)[22,25,35,36]. Therefore, an agent would be able to generate 1,200 synthetic observations of friendships in one minute, equivalent to replaying each friendship approximately 35 times. In contrast, an agent lacking a replay mechanism would only be able to learn from direct experience, and would be limited to the six observations of each friendship from the learning task. We note that 'one minute' of replay does not imply one contiguous minute of uninterrupted replay, but rather one cumulative minute of replay that could be distributed throughout a longer period of time.

Consistent with the hypothesis that replay helps with consolidation, simulation results reveal that even one minute of replay is sufficient to dramatically improve navigation accuracy, compared to an agent that only learns from direct experience (Figure 3C; Methods). However, the simulation also suggests that one minute of replay is nearly as good as one *hour* for consolidating representations (Figure 3C). Therefore, replay-as-consolidation may help to explain how people initially achieve above-chance navigation performance, but it is unlikely to explain navigation improvement after overnight rest. Instead, consistent with an abstraction hypothesis, the simulation reveals that larger values of γ are associated with more accurate navigation decisions, especially for longer distances (Figure 3C).

Given these simulation results, we would expect to empirically observe that human navigation behaviors are better-characterized by larger values of γ after overnight rest. The posterior predictive check again reveals that the SR is able to accurately recapitulate human behavior before and after overnight rest (Figure 3D). As hypothesized, results reveal a significant group-level increase in γ (Day 1 median γ = 0.51, Day 2 median γ = 0.66, 95% CI difference in medians $[0.009, \infty]$, one-tailed p = .023; Figure 3E). To further verify that individual-level changes in estimated γ are associated with greater navigation accuracy, we used Spearman rank correlation to test whether changes in estimated γ track changes in accuracy for shorter- and longer-range navigation problems. Results reveal that 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 3F), but not for the shorter-range problems (distance-2 ρ = -0.18, one-tailed p = .960; Figure 3F). These results are therefore consistent with the proposal that replay affords greater abstraction of a cognitive map that privileges longer-range navigation problems.

As before, we tested the alternative hypothesis that subjects' behaviors were better-described by a model of planning over an internal model, implemented as Breadth-First Search. Results reveal that on both days, the SR model outperforms the BFS model, as well as a hybrid BFS-SR model (all PXP > 0.99).

Offline gains in social navigation rely on cached structural knowledge

Why is longer-range social navigation particularly improved by building more abstract SRs? The longest navigation problems in our studies span the two communities within the network, which necessitates that information flows through an information broker connecting the communities (Figure 1D). Acquiring knowledge about possible routes passing through the broker is therefore critical for solving the longer-range navigation problems, and this knowledge becomes especially useful when information traverses across the communities. In theory, SRs built from greater values of γ should incorporate more multistep connections between the broker and other network members, thus leveraging the broker's fundamental role in information flow. Our simulations reveal that as SRs become more abstract (with larger values of γ), information about multistep connections with the broker is cached (Figure 4A), which helps explain the observed improvement of our subjects in navigating the longest-range navigation problems. In other words, our simulations reveal *how* abstract cognitive maps extract important structural knowledge from more granular knowledge about individual friendships.

If subjects are indeed relying on cached structural knowledge to improve social navigation that spans communities, their performance should be sensitive to structural changes involving the broker, such as the broker breaking off a friendship. An agent relying on cached structural knowledge about the broker's connections should continue to make choices as if no relationship has ruptured, since the agent will either need additional experience to learn about the network change or will need to re-cache the modified multistep relationships through replay in order to correctly navigate the modified network. In contrast, an agent employing model-based planning would be able to rapidly incorporate that change into its internal model and alter its navigation behavior accordingly. To test these alternative models in our subjects, we administered a transition reevaluation procedure on Day 2 in studies 2-3[15,16] (Figure 1F). In a final task, subjects were informed that two people had fallen out and were no longer friends, and that two other people had newly become friends. We engineered these changes such that the critical bridge (i.e., including the broker) between the two communities was severed and formed elsewhere, while no other structural changes were made to either of the two communities (Figure 1D; Figure 1E).

Results reveal that these structural changes to the network were sufficient to abolish subjects' improvements in navigation accuracy after overnight rest (all Ps < .01; Figure 4B, providing strong evidence for the caching of structural knowledge, and evidence against an account of model-based planning. This decrease in navigation accuracy is particularly noteworthy given that subjects had, just thirty minutes prior, exhibited evidence of improved navigation following overnight rest. Indeed, after transition reevaluation, subjects' navigation accuracy was statistically indistinguishable from their initial performance before overnight rest (Figure 4B).

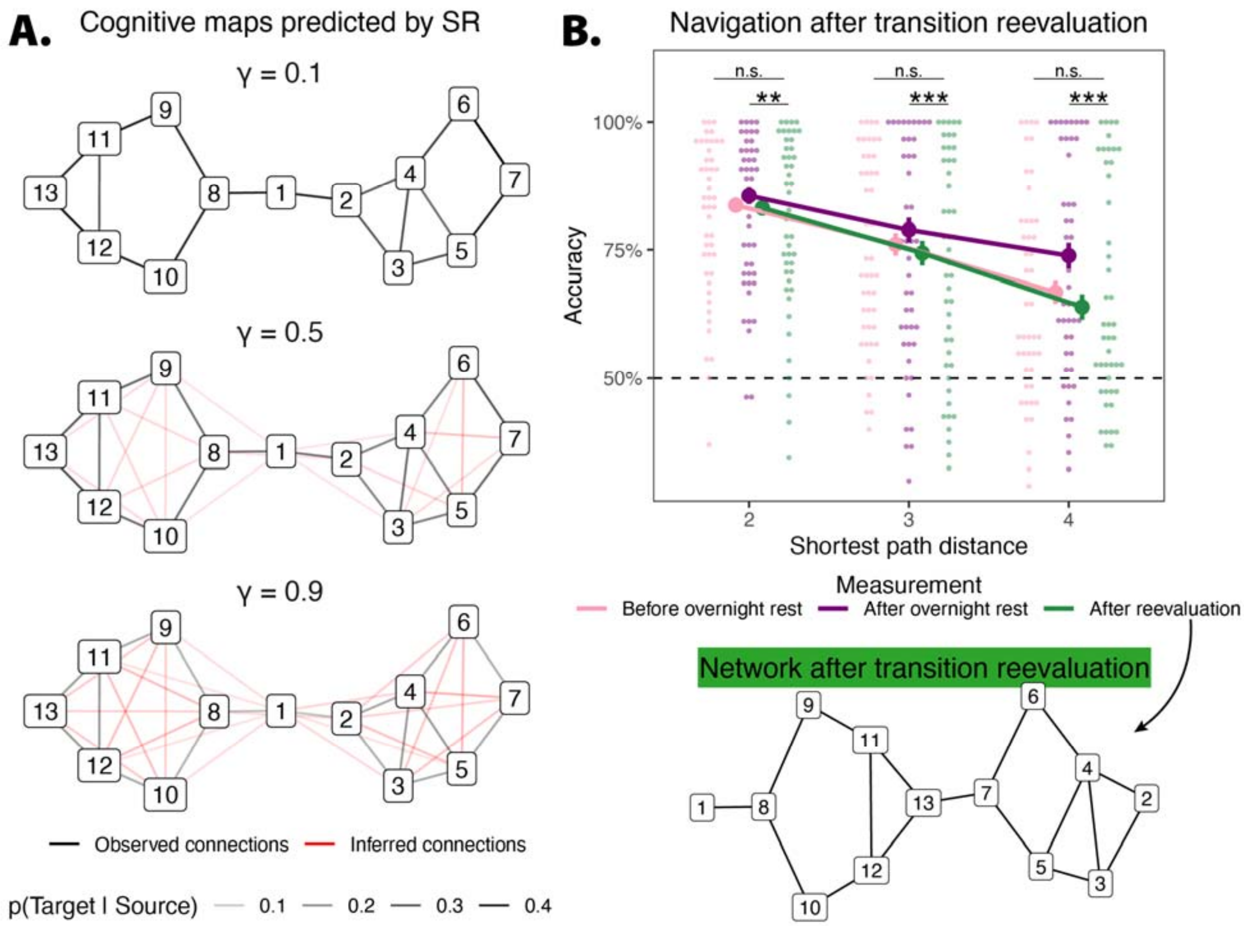

Figure 4 | Evidence of cached structural knowledge. A. *Cognitive maps predicted by the Successor Representation model at different levels of abstraction (γ). When abstraction is low ($\gamma = 0.1$), the cognitive map simply reflects the true network structure. True friendships are color-coded in black. As abstraction increases ($\gamma = 0.5$), the cognitive map begins to reflect inferences about community structure, color-coded in red. With greater abstraction ($\gamma = 0.9$), the cognitive map emphasizes connections to the 'bridging' nodes connecting the two communities, namely nodes 1, 2, and 8. For visual clarity, links have been thresholded at p(Target | Source) > 0.05.* ***B.*** *After being informed of changes in friendship ('transition reevaluation' in which the link between the bridging nodes 1 and 2 was severed), human subjects' navigation accuracy significantly decreased, relative to the improved navigation they had exhibited earlier that same day (i.e., after overnight rest).*

Discussion.

In a set of now-classic studies, Stanley Milgram asked subjects in Nebraska to forward a letter to a target individual they did not know. Subjects were only told the person's name and that they lived in Boston. The job was to mail the letter to someone who could, in turn, forward the letter closer towards the target. Remarkably, of the letters that eventually reached their target, the source and target were only separated by about six degrees[1]. Milgram's study illustrates the fundamental challenge of social navigation: human networks are vast yet densely-connected, meaning that a variety of things—gossip, ideas, norms, disease, and more—are susceptible to being amplified and spread by social networks. To navigate this web of relationships, people must anticipate how information flows, which requires understanding how people are connected. Although this is an inherently difficult problem, Milgram's result suggests that people are surprisingly capable of navigating social networks, even if they lack full knowledge of how people are connected within them. Yet, despite decades of active interest, we still know virtually nothing about the cognitive mechanisms that enable people to solve social navigation problems.

Here, we provide a new experimental framework for closing fundamental gaps in our mechanistic understanding of social navigation. We find that people are proficient at solving social navigation problems requiring inference about how information spreads through a network. Indeed, people can accomplish above-chance navigation accuracy immediately after learning about a novel network, even for longer-range problems that require integrating knowledge over long chains of relationships (friends-of-friends-of-friends-of-friends). Overnight rest further improves social navigation accuracy, and has an especially pronounced effect for performance on problems involving longer-range relationships spanning different communities. Drawing inspiration from decades of research on spatial navigation in rodents and humans, we propose both a representational format enabling information flow to be tracked in the human mind, and the cognitive mechanisms for building these complex mental representations.

First, successful navigation through a network is aided by representing it as an abstract cognitive map encoding not only direct, one-step friendships, but also integrating over indirect, multistep connections like being a friend-of-a-friend. These abstract mental representations can be represented using formats like the Successor Representation and learned using algorithms which can extrapolate multistep relationships from disjointed, pairwise observations of friendship. People's use of multistep abstraction allows them to build more holistic representations of how people in the network are connected to one another, and suggests that abstraction is the lynchpin of how social navigation problems are solved[2,3].

Second, the brain further refines these cognitive maps of social networks during overnight rest using a replay-like mechanism that efficiently reuses experiences from prior learning to generate new, synthetic learning observations. This account is consistent with research showing that animals not only replay prior experiences[22,28,37], but that they also generate entirely new, synthetic 'walks' through the environment[27,38]. Moreover, the fact that we observe the greatest boost in navigational improvement for long-range problems accords with past findings demonstrating that sleep privileges memory abstraction[29-33]. Indeed, we find that extended periods of rest, such as overnight sleep, appear to be involved in building the kinds of highly abstract mental representations that reveal a social network's deeper structure, such as the existence of communities and the individuals that bridge them.

Our studies lay the groundwork for addressing several important questions in future work. Here, we highlight just a few of many promising directions. We establish that a replay-like mechanism is needed to explain how sleep improves navigation performance, but a full account of such a mechanism would require characterizing the content of replay. Past neurobiological findings strongly suggest that replay sequences consist of items that were experienced close together in time (e.g., adjacent locations in a maze that are part of the same path). However, it remains unknown whether this holds true in the context of social networks, where an individual's observations of social interactions may be sequentially or temporally disjointed.

A related question revolves around the neural instantiation of multistep abstraction. Although we leverage the Successor Representation (SR) in this work, we note that multistep abstraction is a much more general representational strategy that could be implemented using many mechanisms with varying degrees of biological plausibility[8,18,19]. Despite the SR depending heavily on the temporal dynamics of experience[13,39], multistep abstraction appears to describe how people represent social networks even when observations of social interaction are temporally disjointed[3]. It is therefore possible that the SR successfully describes social network representation because it is a useful method for discovering structure[14], rather than being a faithful model of neural computation. A particularly intriguing possibility is that the brain may encode components of a network's structure (i.e., basis sets) that afford greater flexibility in assembling useful representations when navigating a variety of social environments[8].

In summary, people track information flow in complex social networks by learning abstract cognitive maps caching knowledge about long-range connections. Our results provide mechanistic insights into how these abstract cognitive maps are learned, and how they are transformed offline by a replay-like mechanism.

METHODS

Subjects.

In study 1, we recruited N = 50 subjects (34 female, 15 male, one nonbinary; mean age = 20.6 years old, SD = 2.81). In study 2, we recruited N = 50 subjects; one subject's demographics were never recorded due to experimenter error. Of subjects whose demographics are known, 31 were female, 17 male, and one nonbinary; the mean age was 23.1 years old, SD = 4.63. In study 3, we recruited N = 50 subjects, but lost four datapoints due to experimenter error, leaving a final sample size of N = 46 (30 female, 16 male; mean age = 23.0 years old, SD = 4.46). All subjects received $10/hour as monetary compensation for their first study session. For the second study session, subjects in study 2 were paid $15, and subjects in study 3 were paid $20. Subjects in studies 2-3 could earn additional cash bonuses of up to $5 depending on how accurately they solved social navigation problems. All study procedures were conducted in a manner approved by the Brown University Institutional Review Board, and all subjects provided informed consent.

Overview.

In study 1, a one-day study (Figure 1F), subjects first learned about a novel social network (Figure 1A), completed a memory test (Figure 1B), and then were tasked with solving social navigation problems (Figure 1C) about the network they had just learned about (Figure 1D). Details about each procedure are provided in subsequent sections.

Study 2 was a two-day study (Figure 1F), where Day 1 was identical to study 1. On Day 2, subjects returned to the lab 24 hours after their first session. In this second session, subjects completed the same memory test and social navigation task as they had in Day 1, then completed the social navigation task a third time after being informed about changes in network members' friendships (Figure 1E).

Study 3 was a two-day study that was nearly identical to study 2, with one key modification. To test the hypothesis that brief awake rest was sufficient to improve navigation accuracy, we added a 15-minute rest period at the end of Day 1, after which subjects completed the memory test and social navigation task again (Figure 1F).

Learning task.

Subjects were required to learn the friendships within an artificial social network. To familiarize subjects with the 13 network members, the task first presented a screen presenting all network members' faces and names, which subjects could examine for as much time as they liked. Afterwards, subjects learned about the friendships between these 13 network members from a computerized 'flashcard' game (Figure 1A). On each trial, subjects were shown one 'Target' network member, and were required to find all of the Target's friends amongst the remaining twelve cards, which were initially displayed face-down. Subjects responded by clicking on face-down cards. Cards flipped face-up and were outlined in green when subjects made correct responses; incorrect responses were indicated by the card remaining face-down and being outlined in red (Figure 1A). Once all of the Target's friends were identified, subjects were given three seconds to review the Target's friends before the task moved on to the next Target.

All network members were presented as Targets, and subjects cycled through all Targets in a single block of trials before moving to the next block. The spatial mapping of network members' cards remained consistent for the first three blocks, then was randomly shuffled for the last three blocks. This was done to ensure that subjects were truly learning about friendships and not simply spatial locations. Overall, the flashcard learning task took 20-25 minutes to complete. All face stimuli were drawn from the Chicago Face Database[40].

Memory test.

Subjects completed a memory test immediately after the learning task (Figure 1B). Each trial presented a Target network member at the top of the screen, and all remaining network members were shown below in two rows of six photographs. Subjects responded by clicking on network members they believed to be the Target's friend. No feedback was ever presented. All responses were self-paced, and the task took 5-10 minutes to complete.

Social navigation task.

On each trial of the 'message-passing' task, subjects chose between two Sources to pass a message to a given Target (Figure 1C). Subjects were explicitly informed that, depending on their choices, the message could be delivered efficiently, inefficiently, or not at all. All primary analyses were performed on trials where there was one unambiguously correct answer (based on shortest path distance). No feedback was ever presented. All responses were self-paced, and the task typically took 30-45 minutes to complete on Day 1 (Figure 1F). This procedure was identical across all three studies.

To test whether a brief period of awake rest was sufficient to improve navigation accuracy, subjects in study 3 completed the initial message-passing task on Day 1, rested for about 15 minutes, then completed the same navigation task again (Figure 1F). During the rest period, subjects completed a task that was designed to keep the social network salient in subjects' minds, while being easy enough that subjects were actually able to rest. For the vast majority of the rest period, subjects were shown a fixation cross on a blank screen. Sporadically, the fixation cross was replaced with a photograph of a network member for 1.5 seconds. Across the entire 15-minute period, 140 photographs were displayed at random, and 50% of them were presented upside-down. The only task was to press a button when a photograph appeared upside-down.

Finally, to test for improvements in navigation performance, subjects completed the same message-passing task on Day 2 in studies 2-3 (Figure 1F). Afterwards, we administered a transition reevaluation procedure to test how alterations in network structure would impact navigation accuracy (Figure 1F). Specifically, we instructed subjects that two individuals who had previously been friends were no longer friends (Figure 1D; Figure 1E), and that a new friendship had been formed between two other network members (Figure 1E). We designed these changes specifically to break a critical bridge between two communities, and create a new bridge elsewhere. These changes in friendship invalidated the longer-range relationships cached by the SR, requiring subjects to quickly adapt to maintain high navigation accuracy. Subjects were not explicitly informed that these changes in friendship fundamentally altered the network's structure.

Behavioral analysis.

We used the R package *glmmTMB* to estimate mixed-effects logistic regression models for the memory and social navigation tasks. Whenever appropriate, we pooled data across the three studies to maximize statistical power and regularize fixed-effects estimates. To account for non-independent observations, the models included random intercepts for each subject, as well as a random intercept for each study. Model-specific random slopes are detailed below.

To test memory accuracy, we pooled data from studies 2-3 and estimated a model where study session (i.e., Day 1 vs Day 2) was both a fixed-effects predictor and a random slope.

In the social navigation task, path distance was defined as the graph distance between the correct Source and Target after removing the ‘letter-writer’ from the network (i.e., because the Source could pass the letter back to the sender). In total, subjects completed a total of 159 trials. Of these, 14 trials presented two Sources that had the same shortest path distance (i.e., both answers were correct), 27 trials required online reevaluation of path distance (i.e., because the shortest possible path would have required sending the letter back to the sender), and 3 trials required both online reevaluation and had the same shortest path distance. Our main analyses focused on the remaining subset of 115 trials where there was an unambiguously correct answer.

We tested two primary hypotheses and two secondary hypotheses. First, we tested whether subjects were able to solve navigation problems with above-chance accuracy on Day 1, immediately after learning about the social network (Figure 1F). As this procedure occurred in all three studies, we pooled data from all studies together. We included predictors for shortest path distance, which we coded as a categorical variable, and additionally estimated random slopes for shortest path distance per subject. To test whether accuracy was above chance at each distance (i.e., shortest paths of 2, 3, and 4), we iteratively re-parameterized the model by making each distance the reference category. Second, we tested whether subjects’ navigation accuracy improved on Day 2 after overnight rest (Figure 1F), and therefore pooled data from studies 2-3. This model included fixed-effects predictors for shortest path distance, the study session (i.e., Day 1 vs Day 2), and their interaction. The model also included per-subject random slopes for shortest path distance and study session. Distance-conditional changes in navigation accuracy were estimated using the same re-parameterization strategy.

Third, we tested the secondary hypothesis that 15 minutes of awake rest would improve navigation accuracy. This model was functionally identical to the two-day model, except that it used data from the ‘Day 1’ and ‘Awake Rest’ navigation tasks in study 3. Finally, we tested whether changes in the network (i.e., transition reevaluation) would result in attenuated navigation accuracy using a model that was functionally identical to the two-day model, except that it used data from the ‘Day 2’ and ‘Transition Reevaluation’ navigation tasks in studies 2-3.

Successor representation model.

In its typical use in reinforcement learning, the Successor Representation (SR) encodes the likelihood that an agent starting at state s will find itself in state t after taking some number of steps dictated by the successor horizon γ (i.e., the lookahead $L = \frac{1}{1-\gamma}$). This has a straightforward translation to social navigation problems like the message-passing task, which requires computing the probability that a message given to a particular Source will be passed to the Target in some number of steps. The SR is encoded as the matrix M with dimensions $N \times N$, where N is the number of network members. Once the SR is learned, an agent could estimate the likelihood that a message given to Source s will make it to Target t simply by looking up the value $M_{s,t}$.

Our implementation used a standard delta-rule method to update M (Equation 1). When network members s and t are observed together, this is encoded in the one-hot vector $\mathbf{1}(t)$, which is a vector of length N filled with zeroes except at the index t. The observation diverges from the agent's prior expectation $M(s)$ (we use this notation to refer to the entire row s, as the SR retrieves and updates M in a row-wise manner), and therefore creates a prediction error. The agent then chains together knowledge of s and t's friendships by adding a fractional amount of $M(t)$ to $M(s)$, controlled by the successor horizon γ. This overall prediction error δ then drives the learning update, tempered by the learning rate α, which was always fixed to 0.1[3,15,21]. As friendships are bidirectional in our study, each learning event prompted two updates, one for s and another for t.

$$M(s) \leftarrow M(s) + \alpha\delta,$$
$$\delta = \mathbf{1}(t) + \gamma M(t) - M(s) \# (Eq.1)$$

We modeled an agent's choice between two candidate Sources using a softmax choice rule (Equation 2), such that the agent retrieves the relevant estimates from $M_{s,t}$, then probabilistically chooses the higher-valued option. Choices are made more deterministically as (inverse) temperature $\tau \to \infty$, and more stochastically as $\tau \to 0$. To account for idiosyncratic and irreducible decision noise, the softmax also included lapse rate $\lambda \in [0, 1]$, such that the agent makes completely random choices when $\lambda = 0$. Due to probabilities being small on an absolute scale, we scaled M by the arbitrary constant 100.

$$p(choose\ Source\ s) = \frac{\exp\left(\tau M_{s,t}\right)}{\exp\left(\tau M_{s,t}\right) + \exp\left(\tau M_{s,t}\right)} \times (1-\lambda) + \frac{\lambda}{2} \#(Eq.2)$$

In the simulation of asymptotic performance, we provided the model with 5,000 observations of each friendship. In the simulation of replay, we treated each novel observation as a single learning event. As the learning task consisted of six blocks, each containing learning events for both $s \rightarrow t$ and $t \rightarrow s$, the no-replay SR learned from 12 observations of each friendship.

Maximum-likelihood parameter-fitting was performed using the Nelder-Mead algorithm implemented in R's default optimizer. Parameters were fit independently for each subject, and each model was re-estimated 25 times, keeping only the estimates that best maximized the likelihood.

Model-based planning.

We tested an alternative computational model of breadth-first search (BFS), a normatively optimal method for solving shortest path problems given an internal model. Traditional BFS is guaranteed to find the shortest path between two network members, which would result in perfect accuracy in our task. To give the BFS model the capacity to generate more humanlike behavior, we made three key modifications.

First, we introduced realistic decision noise stemming from the need to search two Sources at the same time. We assumed that the agent would start two separate tree searches, one from each Source. On each iteration, the agent randomly chose to search Source A or B further. When first searching a given Source, the agent would retrieve all of that Source's friends. Subsequent searches would retrieve friends-of-friends, friends-of-friends-of-friends, and so on. We assumed that the agent had the capacity to remember what network members had already been retrieved in each of the two searches, and the agent stopped searching once the Target was discovered. To our knowledge, there is no analytic method for computing choice and 'reaction time' distributions from such a search process, so we simulated how our BFS agent would solve each trial, 5,000 times per trial.

Second, we estimated a 'search threshold' parameter that made an agent more likely to choose randomly during long searches. For example, an agent with a search threshold of six was likely to continue searching for the Target after retrieving a few network members in memory. However, after retrieving six network members, that agent would have a 50% chance of giving up and choosing randomly; the probability of choosing randomly would continue increasing with the length of the search. Third, and finally, we estimated a lapse rate parameter capturing a general tendency to choose randomly, unrelated to the length of any particular search.

We also estimated a hybrid BFS-SR model, which was nearly identical to the BFS model. However, when choosing to give up (i.e., after the search threshold was exceeded), the agent fell back on using cached SR estimates instead of choosing randomly. Additionally, due to concerns about parameter identifiability, no lapse rate was estimated for the hybrid model.

To fit parameters to subjects' behavior, we defined logistic loss as the difference between a subject's choice on a particular trial, and the average simulated choice (from 5,000 iterations) of the BFS agent on that trial. The search threshold parameter was estimated as the value at which a softmax with $\tau = 1$ was indifferent between choosing to complete a search (based on the average length of the BFS search for a given trial) and giving up. Likelihoods were weighted accordingly. For example, if an agent was estimated to be 60% likely to give up during a particularly long search, the BFS prediction contributed 40% to the overall likelihood.

Computational model comparison.

Protected exceedance probabilities provide a formal test of a model's group-level fit compared to other candidate models[41], and were computed using R software written by Matteo Lisi (https://github.com/mattelisi/bmsR).

ACKNOWLEDGEMENTS
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.).

AUTHOR CONTRIBUTIONS

Conceptualization: M-L.V., J.Y.S., A.B., O.F.H. Formal analysis: J.Y.S., M-L.V. Funding acquisition: O.F.H. Investigation: M-L.V. Methodology: M-L.V., J.Y.S., A.B., O.F.H. Supervision: O.F.H., A.B., Writing: J.Y.S., M-L.V., A.B., O.F.H.

DATA AND CODE AVAILABILITY

All data and code needed to reproduce the analyses are available in a publicly-accessible GitHub repository: https://github.com/feldmanhalllab/network-navigation-replay

REFERENCES

1 Travers, J. & Milgram, S. An Experimental Study of the Small World Problem. *Sociometry* **32**, 425-443 (1969). https://doi.org:10.2307/2786545

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

3 Son, J.-Y., Bhandari, A. & FeldmanHall, O. Abstract cognitive maps of social network structure aid adaptive inference. *Proceedings of the National Academy of Sciences* **120**, e2310801120 (2023). https://doi.org:10.1073/pnas.2310801120

4 Tolman, E. C. Cognitive maps in rats and men. *Psychological Review* **55**, 189-208 (1948). https://doi.org:10.1037/h0061626

5 O'Keefe, J. & Nadel, L. *The Hippocampus as a Cognitive Map*. (Oxford: Clarendon Press, 1978).

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

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

8 Behrens, T. E. J. *et al.* What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron* **100**, 490-509 (2018). https://doi.org:10.1016/j.neuron.2018.10.002

9 Constantinescu, A. O., O'Reilly, J. X. & Behrens, T. E. J. Organizing conceptual knowledge in humans with a gridlike code. *Science* **352**, 1464 (2016). https://doi.org:10.1126/science.aaf0941

10 Garvert, M. M., Dolan, R. J. & Behrens, T. E. J. A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *eLife* **6**, e17086 (2017). https://doi.org:10.7554/eLife.17086

11 Tavares, Rita M. *et al.* A Map for Social Navigation in the Human Brain. *Neuron* **87**, 231-243 (2015). https://doi.org:https://doi.org/10.1016/j.neuron.2015.06.011

12 Park, S. A., Miller, D. S., Nili, H., Ranganath, C. & Boorman, E. D. Map Making: Constructing, Combining, and Inferring on Abstract Cognitive Maps. *Neuron* **107**, 1226-1238.e1228 (2020). https://doi.org:10.1016/j.neuron.2020.06.030

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

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

15 Momennejad, I. *et al.* The successor representation in human reinforcement learning. *Nature Human Behaviour* **1**, 680-692 (2017). https://doi.org:10.1038/s41562-017-0180-8

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

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

18 Lynn, C. W. & Bassett, D. S. How humans learn and represent networks. *Proceedings of the National Academy of Sciences* **117**, 29407 (2020). https://doi.org:10.1073/pnas.1912328117

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

20 Momennejad, I. & Howard, M. W. Predicting the Future with Multi-scale Successor Representations. *bioRxiv*, 449470 (2018). https://doi.org:10.1101/449470

21 Pudhiyidath, A. *et al.* Representations of Temporal Community Structure in Hippocampus and Precuneus Predict Inductive Reasoning Decisions. *Journal of Cognitive Neuroscience*, 1-25 (2022). https://doi.org:10.1162/jocn_a_01864

22 Foster, D. J. Replay Comes of Age. *Annual Review of Neuroscience* **40**, 581-602 (2017). https://doi.org:10.1146/annurev-neuro-072116-031538

23 Schapiro, A. C., McDevitt, E. A., Rogers, T. T., Mednick, S. C. & Norman, K. A. Human hippocampal replay during rest prioritizes weakly learned information and predicts memory performance. *Nature Communications* **9**, 3920 (2018). https://doi.org:10.1038/s41467-018-06213-1

24 Sun, W., Advani, M., Spruston, N., Saxe, A. & Fitzgerald, J. E. Organizing memories for generalization in complementary learning systems. *Nature Neuroscience* **26**, 1438-1448 (2023). https://doi.org:10.1038/s41593-023-01382-9

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

26 Foster, D. J. & Wilson, M. A. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**, 680-683 (2006). https://doi.org:10.1038/nature04587

27 Igata, H., Ikegaya, Y. & Sasaki, T. Prioritized experience replays on a hippocampal predictive map for learning. *Proceedings of the National Academy of Sciences* **118**, e2011266118 (2021). https://doi.org:10.1073/pnas.2011266118

28 Zhenglong, Z., Michael, J. K. & Anna, C. S. Replay as context-driven memory reactivation. *bioRxiv*, 2023.2003.2022.533833 (2023). https://doi.org:10.1101/2023.03.22.533833

29 Ellenbogen, J. M., Hu, P. T., Payne, J. D., Titone, D. & Walker, M. P. Human relational memory requires time and sleep. *Proceedings of the National Academy of Sciences* **104**, 7723-7728 (2007). https://doi.org:10.1073/pnas.0700094104

30 Lewis, P. A. & Durrant, S. J. Overlapping memory replay during sleep builds cognitive schemata. *Trends in Cognitive Sciences* **15**, 343-351 (2011). https://doi.org:10.1016/j.tics.2011.06.004

31 Lutz, N. D., Diekelmann, S., Hinse-Stern, P., Born, J. & Rauss, K. Sleep Supports the Slow Abstraction of Gist from Visual Perceptual Memories. *Scientific Reports* **7**, 42950 (2017). https://doi.org:10.1038/srep42950

32 Feld, G. B., Bernard, M., Rawson, A. B. & Spiers, H. J. Sleep targets highly connected global and local nodes to aid consolidation of learned graph networks. *Scientific Reports* **12**, 15086 (2022). https://doi.org:10.1038/s41598-022-17747-2

33 Klinzing, J. G., Niethard, N. & Born, J. Mechanisms of systems memory consolidation during sleep. *Nature Neuroscience* **22**, 1598-1610 (2019). https://doi.org:10.1038/s41593-019-0467-3

34 Correa, C. G., Ho, M. K., Callaway, F., Daw, N. D. & Griffiths, T. L. Humans decompose tasks by trading off utility and computational cost. *PLOS Computational Biology* **19**, e1011087 (2023). https://doi.org:10.1371/journal.pcbi.1011087

35 Kurth-Nelson, Z., Economides, M., Dolan, Raymond J. & Dayan, P. Fast Sequences of Non-spatial State Representations in Humans. *Neuron* **91**, 194-204 (2016). https://doi.org:https://doi.org/10.1016/j.neuron.2016.05.028

36 Schuck, N. W. & Niv, Y. Sequential replay of nonspatial task states in the human hippocampus. *Science* **364**, eaaw5181 (2019). https://doi.org:10.1126/science.aaw5181

37 Jadhav, S. P., Kemere, C., German, P. W. & Frank, L. M. Awake Hippocampal Sharp-Wave Ripples Support Spatial Memory. *Science* **336**, 1454-1458 (2012). https://doi.org:10.1126/science.1217230

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

39 Gershman, S. J., Moore, C. D., Todd, M. T., Norman, K. A. & Sederberg, P. B. The Successor Representation and Temporal Context. *Neural Computation* **24**, 1553-1568 (2012). https://doi.org:10.1162/NECO_a_00282

40 Ma, D. S., Correll, J. & Wittenbrink, B. The Chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods* **47**, 1122-1135 (2015). https://doi.org:10.3758/s13428-014-0532-5

41 Rigoux, L., Stephan, K. E., Friston, K. J. & Daunizeau, J. Bayesian model selection for group studies — Revisited. *NeuroImage* **84**, 971-985 (2014). https://doi.org:https://doi.org/10.1016/j.neuroimage.2013.08.065