nature human behaviour

Article https://doi.org/10.1038/s41562-024-01990-w

Replay shapes abstract cognitive maps for efficient social navigation

Received: 3 January 2024

Accepted: 16 August 2024

Published online: 19 September 2024

Check for updates

Jae-Young Son[1], Marc-Lluís Vives[2], Apoorva Bhandari[1,4] & Oriel FeldmanHall[1,3,4]

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 first-hand experience with every possible route between individuals. How, then, are people able to accurately track information flow through social networks? Here we find that people immediately 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 while 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 privileges improvements in social navigation accuracy for the longest communication paths that span distinct communities within the network. Together, these findings provide mechanistic insight into the sophisticated mental representations humans use for social navigation.

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 to 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[2]. To navigate this web of relationships, people must anticipate how information flows, which requires understanding how people are connected[3,4]. 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, little is known about the cognitive mechanisms that enable people to solve social navigation problems.

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[5–7]. The format of these spatial maps allows objects to be placed within two-dimensional mental spaces[6,8], affording representation of the longer-range relationships between those objects. Outside of spatial navigation, recent work demonstrates that humans also represent abstract maps of conceptual spaces[9–11], including social traits such as competence and popularity[12,13]. However, 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 candidate representational format for social networks that encodes not only the direct connections between entities (for example, friendships), but also longer-range, multi-step connections (for example, friends of friends)[11,14–20]. This abstracted representation of social networks is related to successor representations (SRs) in reinforcement learning and can be learned using simple, biologically plausible mechanisms.

[1]Department of Cognitive 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. [4]These authors contributed equally: Apoorva Bhandari, Oriel FeldmanHall. e-mail: apoorva_bhandari@brown.edu; oriel.feldmanhall@brown.edu

https://doi.org/10.1038/s41562-024-01990-w

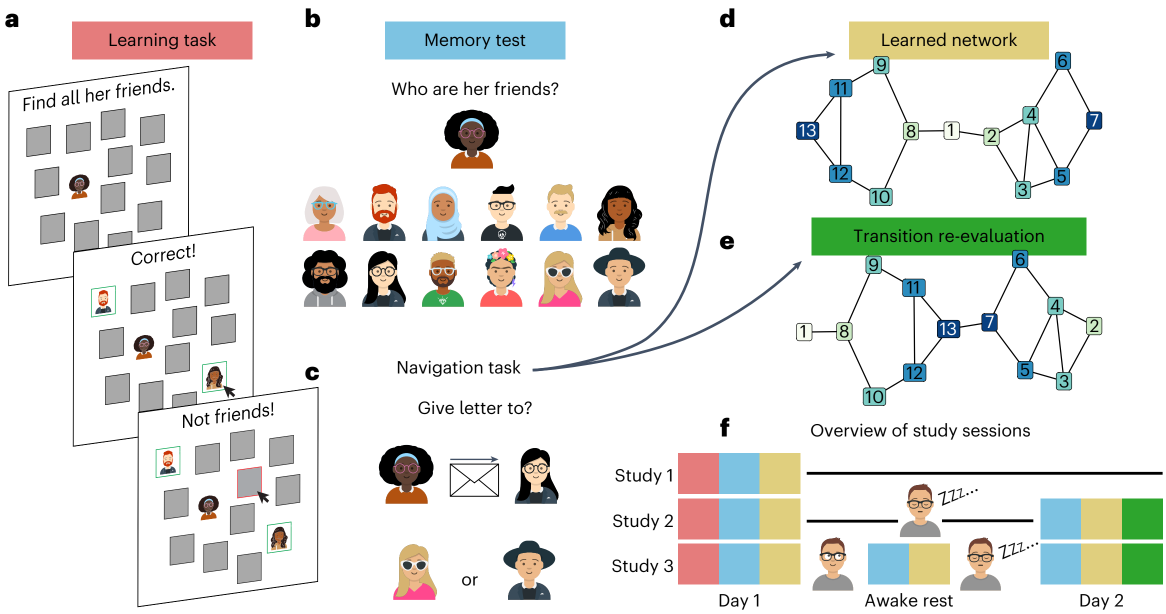

Fig. 1 | Study design. a, The learning task used 'flashcards' to facilitate rapid and accurate learning. When presented with a target network member, subjects were required to find all the target's friends from the face-down cards. When a card corresponding to the target's friend was selected, the card was 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 subjects 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. Subjects were required to indicate which source was the better choice for efficient delivery to the target. **d**, The social network learned by subjects. **e**, In studies 2 and 3, subjects were informed that some of the friendships had been broken and that others had formed. This necessitated rapid re-evaluation of how network members were related to one another. **f**, A schematic illustrating what tasks subjects completed on what days, in which studies. The colour coding corresponds to the task labels in **a**–**e**. Yellow and green indicate completion of the navigation task for the learned and re-evaluated networks, respectively. All avatar icons were generated using getavataaars.com, designed by Pablo Stanley and developed by Fang-Pen Lin.

By adjusting how many steps are integrated over, network representations can be learned at various levels of abstraction[21], where greater abstraction confers rapid inference about distant relations, as well as the existence of network structures, such as communities[4,15,22,23]. This ability to represent longer-range relations is likely to aid 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 and human neuroscience points to an important role of replay, where the brain generalizes from experience to simulate synthetic observations that can drive additional learning, especially prioritizing those that are most critical for adaptive navigation[24–30]. Indeed, it has long been noted that a replay-like mechanism appears necessary to learn sufficiently abstract representations for navigation[15–17]. Offline replay during sleep appears to play an especially important role in generating more abstract representations of the environment[31–35]. As abstraction can help reveal the underlying structure of a given environment and therefore aid longer-range navigation[15,18], it is probable that extended periods of rest, such as overnight sleep, are critical for building the kinds of abstract representations needed for longer-range navigation through social networks.

Therefore, multi-step abstraction not only specifies a useful format for representing the topology of graph structures, such as social networks[4,11,16–18,23], but also provides a natural interface between cognitive maps and replay. Although multi-step 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 multi-step abstraction[4], and it is yet unknown whether, or how, multi-step abstraction supports navigation behaviours.

Here, we test whether humans rely on cognitive maps of social networks for social navigation and if a replay-like mechanism supports more successful social navigation. To assess humans' use of cognitive maps to solve the challenge of social navigation, we created a task where subjects learn about friendships in a social network (Fig. 1a), allowing us to probe how people navigate information flow through a community (Fig. 1c). We then either had subjects take the navigation task immediately, or brought subjects back to the laboratory the next day to test whether navigation accuracy improved after overnight rest (Fig. 1f). Using computational modelling, we characterized the underlying cognitive maps employed by subjects and further tested whether a replay-like mechanism helps to scaffold more successful social navigation.

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 (Fig. 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 friends with the target directly, we classified these as distance-two 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. The two sources presented on each trial were always friends of the message sender to prevent potential confounds, and 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 were replicated across all studies; Supplementary Table 3), subjects completed this navigation task shortly after learning friendships in a novel social network. (Fig. 1a). Subjects never observed the whole network and were provided no direct information about multi-step, longer-range connections between network members (for example, 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 two, regression coefficient (b) = 1.68, z-score (z) = 14.30, 95% confidence interval (CI) 1.45 to 1.91, $P < 0.001$), but also for the longer-range problems (70% accuracy at distance three, $b = 1.06$, $z = 9.34$, 95% CI 0.84 to 1.28, $P < 0.001$ and 63% accuracy at distance four, $b = 0.66$, $z = 6.05$, 95% CI 0.44 to 0.87, $P < 0.001$, all results from mixed-effects logistic regression; Fig. 2a and Methods). 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.

Computational models of social navigation

We consider two classes of decision-making strategies that an agent could employ to flexibly solve novel social navigation problems: model-based planning and caching abstracted representations of multi-step relations. A normatively optimal agent would represent an internal model of 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 practise, online navigation of this kind is time-consuming and computationally costly[17], but can be made tractable in small networks using search algorithms, such as breadth-first search (BFS) that have previously been studied as cognitive 'path-finding' models[36,37]. Here, we test the following planning-based models: (1) BFS-forwards (BFS-F), in which the agent performs two forwards searches from each of the two candidate sources, choosing the source where the target is first found, (2) BFS-backwards (BFS-B), in which the agent performs a single search starting from the target, choosing the first source that is found, and (3) an ideal observer, which computes the shortest path distance between each source and target, choosing the source that is closer to the target. To make these planning models more psychologically plausible, our implementations included parameters that captured the human tendency to give up and choose randomly during long searches, as well as decision noise when choosing between two options (Methods).

Alternatively, an agent could navigate more efficiently by caching (that is, pre-computing) relevant knowledge. In the context of social navigation, it would be particularly useful to cache knowledge of individuals' longer-range, multi-step connections (for example, friends of friends). Recent work in cognitive neuroscience points to the SR as a useful format for encoding such cached knowledge[16–18], including cognitive maps of social networks[4]. 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- versus longer-range connections. As $\gamma \to 0$, the agent represents shorter-range relations, such that the SR only encodes one-step relations (that is, direct friendships) when $\gamma = 0$. As $\gamma \to 1$, the agent integrates over longer-range connections (for example, friends of friends of friends and so on). 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 β, controlling how noisily the agent chooses.

Simulation results reveal, a priori, that multi-step 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 (Fig. 2b and Methods). These results therefore confirm that human subjects' social navigation decisions could in principle be supported by a cognitive map of multi-step relationships, where representation of longer-range relations is supported by greater abstraction (that is, larger γ; Fig. 2c).

We next fitted this model to subjects' behaviours before rest to test whether multi-step abstraction quantitatively outperforms model-based planning in explaining human behaviour in the social navigation task (Methods). We used protected exceedance probability (PXP) to formally test the probability that one model provided a superior fit over all other models under consideration[38]. The results revealed that the SR indeed provided a better group-level fit to the data than all online planning models (all PXP > 0.97; Fig. 2d and Methods), mirrored in the Akaike weights[39] and the proportion of subjects best fitted by the SR (Fig. 2d). Therefore, a formal comparison of computational models suggests that human behaviour on the message-passing task is best explained by the use of a cognitive map containing cached knowledge of abstract, multi-step relations. Posterior predictive checks further confirmed that the planning-based models systematically mischaracterize human subjects' navigation behaviours, while the SR model is largely successful in recapitulating human behaviour (Fig. 2e).

Finally, we tested whether each model was able to predict human behaviour in a held-out subset of navigation problems (that is, trials that were not used to fit model parameters) before rest. 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 planning models do not systematically favour one source over the other in these trials (Fig. 2f), humans demonstrate preferences for sources that have multiple (relatively) short paths to the target, which is mirrored by the SR model predictions (Fig. 2f and Supplementary Fig. 13).

Social navigation improves with overnight rest

To test whether a replay-like mechanism might result in improved social navigation after an extended period of overnight rest that includes sleep, subjects in studies 2 and 3 ($N = 96$) completed a 2 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 (distance-two accuracy 81% before rest, 82% after overnight rest, $b = 0.23$, $z = 2.91$, 95% CI 0.08 to 0.39, $P = 0.004$; distance-three accuracy 73% before rest, 75% after overnight rest, $b = 0.26$, $z = 2.99$, 95% CI 0.09 to 0.44, $P = 0.003$; distance-four accuracy 65% before rest, 71% after overnight rest, $b = 0.43$, $z = 5.17$, 95% CI 0.27 to 0.59, $P < 0.001$, all results from mixed-effects logistic regression; Fig. 3a). This improvement was particularly pronounced for the longest-range distance-four problems compared with the accuracy improvement for distance-two problems ($b = 0.20$, $z = 2.36$, 95% CI 0.03 to 0.36, $P = 0.018$).

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 min at the end of the standard day 1 procedure and before overnight rest (Fig. 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 $P > 0.1$; Fig. 3b), suggesting that a longer period of rest (or possibly sleep) may be needed to produce significant improvements in navigation.

A computational model of replay

In the theoretical framework of the SR, replay is a natural mechanism for explaining how overnight rest improves social navigation[15–17]. 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

 https://doi.org/10.1038/s41562-024-01990-w

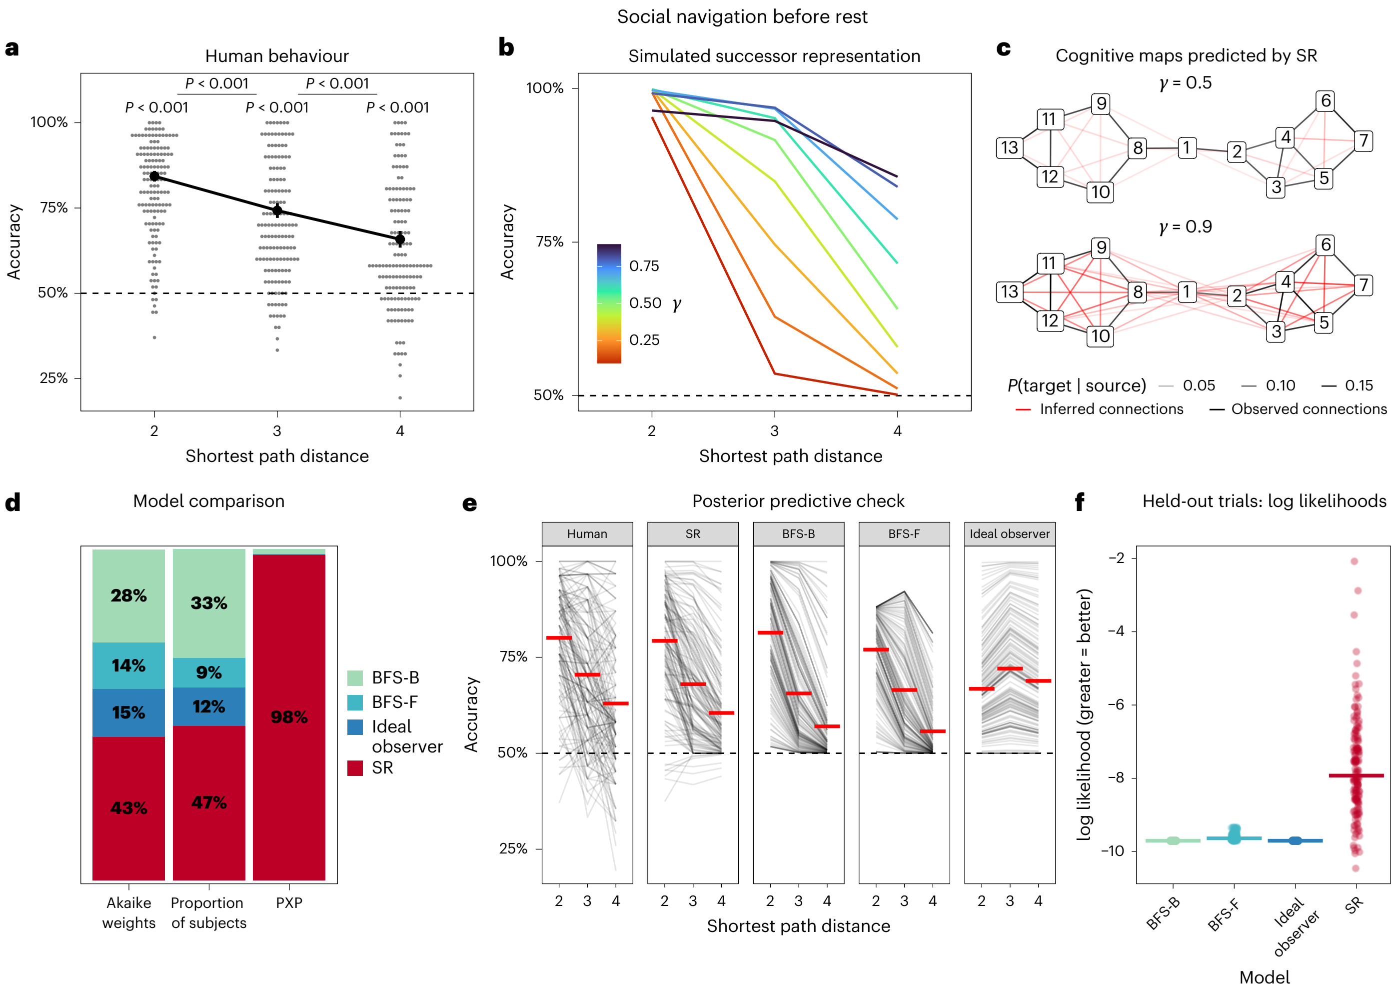

Fig. 2 | Evidence of social navigation. a, Shortly after learning about friendships in a novel social network, people are able to solve social navigation problems with above-chance accuracy. The trend lines reflect the estimated means from a mixed-effects logistic regression model, the error bars reflect estimated standard errors ($N = 146$) and P values are from two-sided tests. **b**, Simulated navigation behaviour from the SR model, which learns multi-step relations between network members over a horizon controlled by γ. The softmax inverse temperature is fixed at $\beta = 100$ for visualization. **c**, A cognitive map learned with greater γ allows efficient representation of longer-range relationships. True friendships are colour coded in black. With increasing abstraction, the cognitive map begins to reflect inferences about community structure and connections to the 'bridging' nodes connecting the two communities. For visualization, links have been thresholded at P(target | source) > 0.025. **d**, The SR is the best-fitting model in terms of Akaike weights, the proportion of subjects best-fit and PXP. **e**, The simulations of computational model behaviours are based on the parameters estimated from human subject behaviours. The bars reflect group-level means. **f**, In a subset of 'held-out' navigation problems (that is, not used to fit parameters), 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. On these held-out trials, the SR has greater out-of-sample likelihoods than the planning models, demonstrating that it is better able to explain human subjects' preferences. The bars reflect group-level medians ($N = 146$). The dashed horizontal lines reflect chance-level accuracy.

hypotheses of how SR 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[15,25]. 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 multi-step relations (that is, with larger γ)[15,16,21]. Intuitively, SR replay is likely to be less comprehensive during brief periods of awake rest compared with extended periods of overnight rest, and it is therefore possible that overnight rest helps to stitch knowledge of pairwise relationships into representations of longer-range multi-step 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 SR replay. For simplicity, we assumed that the agent replayed all friendships in the network in one 'iteration' before replaying another iteration (that is, 17 undirected relations and 34 directed). Past empirical studies of neural replay have found that it takes approximately 50 ms to replay a single transition between two states (that is, a friendship dyad)[24,27,40,41], such that an SR agent could replay 50 iterations in the span of approximately 90 s. In contrast, an agent lacking an SR 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, in translating iterations of SR replay to absolute time, we do not imply contiguous, uninterrupted replay, but rather cumulative SR replay that could be distributed throughout a longer period of absolute time.

Consistent with the hypothesis that SR replay helps with consolidation, simulation results reveal that even a small amount of SR replay is sufficient to dramatically improve navigation accuracy compared with an agent that only learns from direct experience (Fig. 3c and Methods). However, the simulation also suggests that the benefits gained from consolidation rapidly asymptote (Fig. 3c). Therefore,

https://doi.org/10.1038/s41562-024-01990-w

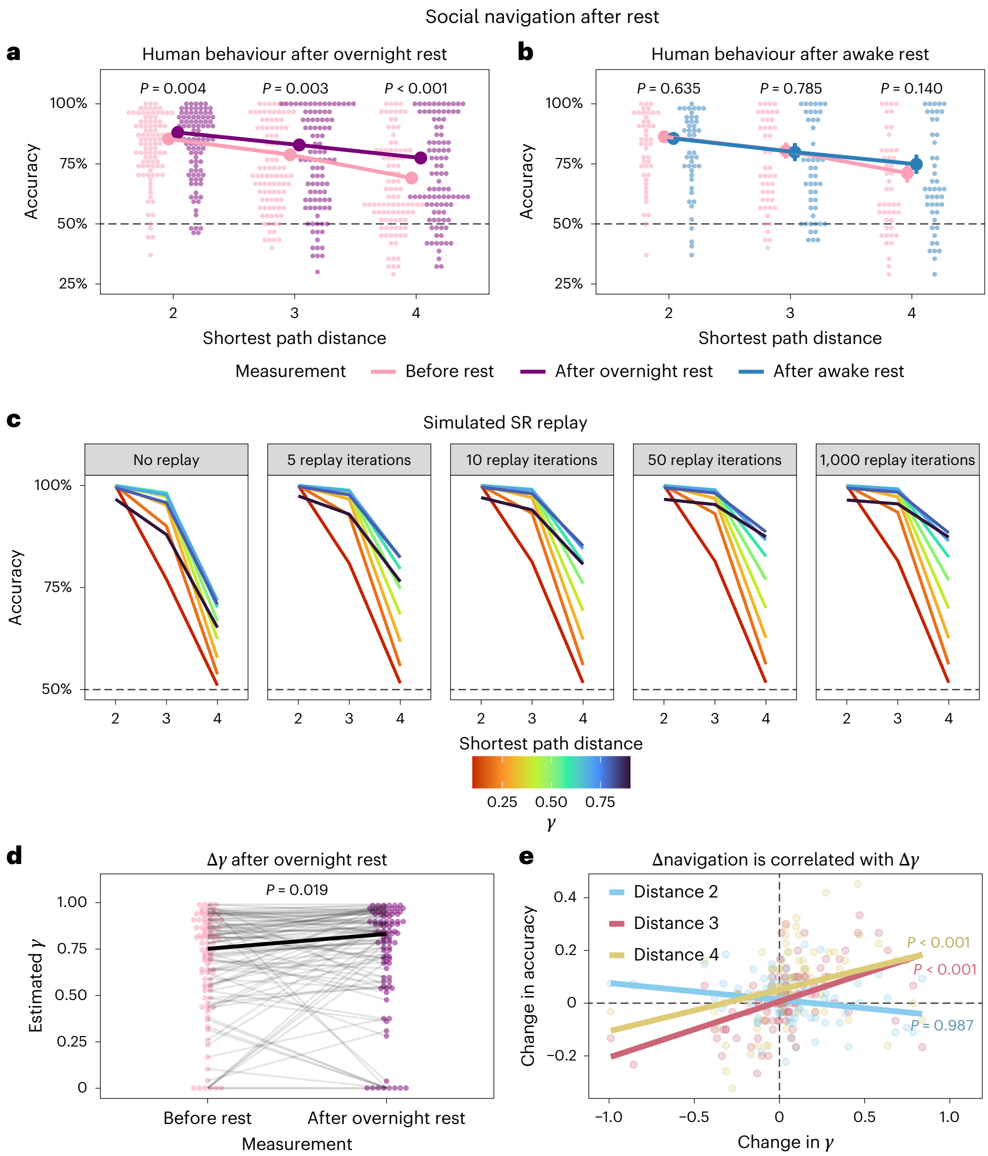

Fig. 3 | Evidence for a replay-like mechanism. a, After overnight rest, humans became more accurate at solving social navigation problems. Increases in navigation accuracy were particularly pronounced for problems requiring knowledge of longer-range connections. **b**, Brief awake rest did not significantly improve navigation. In **a** and **b**, the trend lines reflect the estimated means from a mixed-effects logistic regression model, the error bars reflect estimated standard errors (N = 96 and 46, respectively) and P values are from two-sided tests. **c**, Simulations suggest that even a small amount of SR replay helps an agent consolidate its representation of friendships in a social network, resulting in more accurate navigation. However, the benefits from consolidation rapidly asymptote, 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 multi-step relations (for example, friends of friends of friends, and so on) and which aid in solving longer-range navigation problems. **d**, Parameter estimates from the SR model fit to subjects' navigation behaviours reveal a group-level increase in γ after overnight rest. The thick line reflects the group-level medians (N = 96), and P values are from a one-sided test. **e**, The 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 results from one-sided Spearman rank correlation, Bonferroni corrected for three comparisons. The dashed horizontal lines reflect chance-level accuracy, except in **e**.

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 (Figs. 2c and 3c).

Given these simulation results and past research showing that a variety of mental representations become more abstract during sleep[31–35,42–46], it may be possible that extended periods of rest enable the brain to replay longer sequences through the network during overnight rest. Empirically, this would be reflected in human navigation behaviours being better characterized by larger values of γ after overnight rest. As hypothesized, the results reveal a significant group-level increase in γ (day 1 median $\gamma = 0.75$, day 2 median $\gamma = 0.83$, 95% CI difference in medians [0.009 to ∞], one-tailed $P = 0.019$; Fig. 3d). 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. The results reveal that increased γ on day 2 was associated with improved navigation accuracy for the longer-range problems (distance-three $\rho = 0.55$, one-tailed $P < 0.001$ and distance-four $\rho = 0.47$, one-tailed $P < 0.001$; Fig. 3e) but not for the shorter-range problems (distance-two $\rho = -0.23$, one-tailed $P = 0.987$; Fig. 3e). 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' behaviours were better described by model-based planning. The results reveal that, on both days, the SR model outperforms all planning models (all PXP > 0.97). Finally, in addition to the formal model comparison favouring the SR over the planning models, the results also reveal that subjects' memory performance significantly decreased after overnight rest ($b = -0.29$, $z = -4.65$, 95% CI −0.41 to −0.17, $P < 0.001$). This is inconsistent with a theory of model-based planning, as decreased memory performance suggests that an individual's internal model gets worse, not better, following overnight rest.

Offline gains in social navigation rely on cached 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 groups (Fig. 2c). 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 multi-step 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 multi-step connections with the broker is cached (Fig. 2c), 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 either will need additional experience to learn about the network change or will need to re-cache the modified multi-step relationships through replay (either offline or 'on-task') 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 behaviour accordingly. To test whether subjects show evidence of using cached representation, we administered a transition re-evaluation procedure on day 2 in studies 2 and 3 (refs. 16,17) (Fig. 1f). In a final task, subjects were informed that two people were no longer friends and that two other people had newly become friends. We engineered these changes such that the critical bridge (that is, 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 (Fig. 1d,e).

The results reveal that these structural changes to the network were sufficient to abolish subjects' improved accuracy for longer-range navigation problems after overnight rest (distance-three $b = 0.20$, $z = 2.33$, 95% CI 0.03 to 0.37, $P = 0.020$ and distance-four $b = 0.45$, $z = 5.51$, 95% CI 0.29 to 0.61, $P < 0.001$), but not shorter-range problems (distance-two $b = 0.14$, $z = 1.77$, 95% CI −0.02 to 0.30, $P = 0.077$; Fig. 4a and Supplementary Table 6), providing evidence that subjects relied on cached structural knowledge to inform longer-range social navigation decisions. This decrease in navigation accuracy is particularly noteworthy given that subjects had, just 30 min prior, exhibited evidence of improved navigation following overnight rest. Indeed, after transition re-evaluation, subjects' navigation accuracy was statistically indistinguishable from their initial performance before overnight rest (Fig. 4b and Supplementary Table 6).

To formally test whether the cached representation was more consistent with multi-step abstraction, rather than model-based planning, we took advantage of the fact that subjects completed two navigation tasks on the same day (that is, before and after transition re-evaluation). This aspect of the study design enabled us to predict post re-evaluation behaviour using the computational model parameters that had previously been estimated from pre re-evaluation behaviour, such that the data used to fit the models' parameters was completely different from the data used to test the models' goodness of fit (Methods). Not only does the SR has the highest out-of-sample likelihood weights (conceptually similar to Akaike weights, but based on raw log likelihood rather than the Akaike Information Criterion (AIC), but it is also the absolute best-fitting model for the highest proportion of subjects, and is favoured over all planning models in a Bayesian model selection analysis (PXP = 0.98; Fig. 4b).

Given empirical evidence that subjects cached representations of multi-step relationships, we used simulation to test how an SR agent might use 'on-task' replay to quickly re-cache the structure of the network after transition re-evaluation[47]. The simulation results reveal that an SR agent that does no updating (that is, continues to use the cached representation it had learned for the original network) is generally unable to solve distance-four problems, with accuracy falling at, or even below, chance levels (Fig. 4c). In contrast, an SR agent that learns from replay exhibits dramatic gains in accuracy after even one iteration of replay (that is, replaying all of the friendships in the re-evaluated network; Fig. 4c). Therefore, the simulation results demonstrate that even a relatively small amount of updating is sufficient for explaining how cached SRs can achieve above-chance navigation accuracy after transition re-evaluation.

Discussion

Stanley Milgram's seminal studies, more than 50 years ago, demonstrated that people are able to efficiently pass messages through a large, complex social network, hinting at a human capacity for representing social relationships in a format that supports social navigation[1]. Here, we provide a new experimental framework for closing fundamental gaps in our mechanistic understanding of how people adaptively navigate social relationships. 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, multi-step connections such as being a friend of a friend. These abstract mental representations can be learned using algorithms that extrapolate multi-step relationships from disjointed, pairwise observations of friendship. People's use of multi-step 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 linchpin of how social navigation problems are solved[3,4].

https://doi.org/10.1038/s41562-024-01990-w

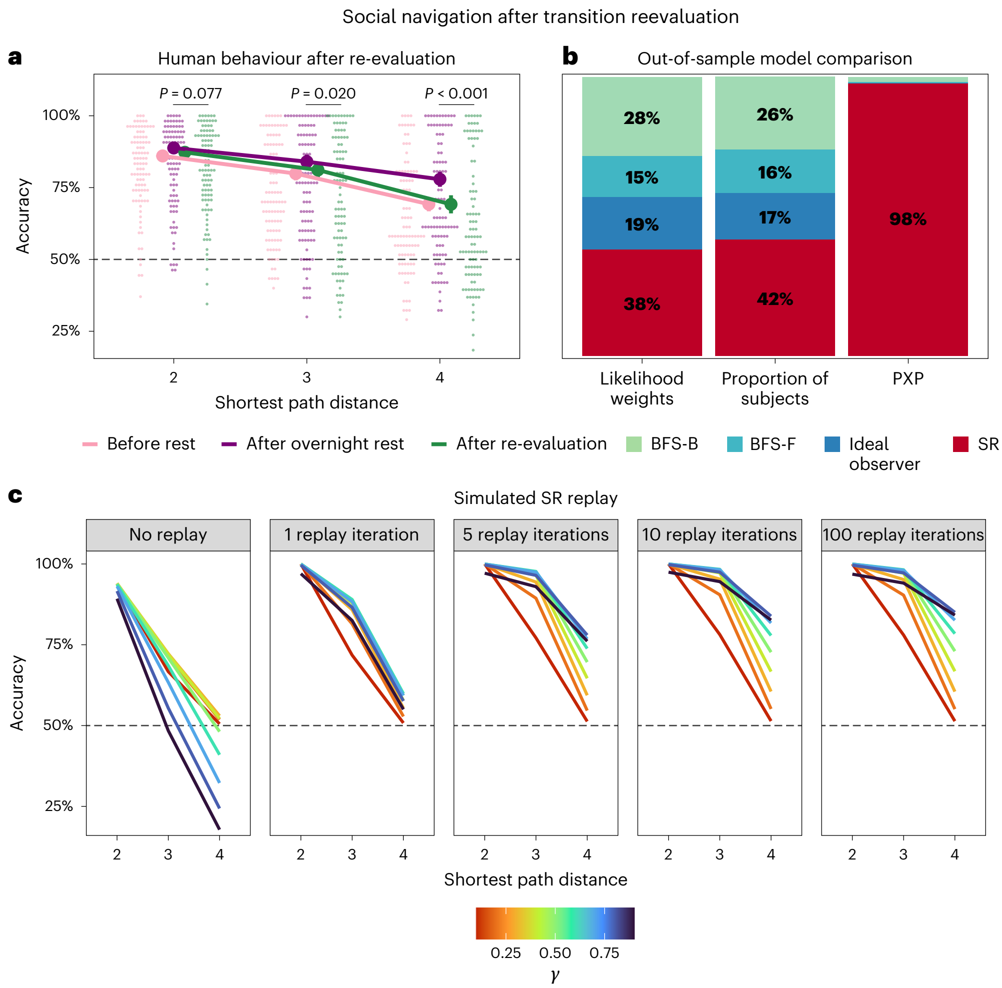

Fig. 4 | Evidence of cached structural knowledge. a, After being informed of changes in friendship ('transition re-evaluation' in which the link between the bridging nodes was severed), human subjects' navigation accuracy significantly decreased, relative to the improved navigation they had exhibited earlier that same day (that is, after overnight rest). The trend lines reflect the estimated means from a mixed-effects logistic regression model, the error bars reflect estimated standard errors ($N = 96$) and P values are from two-sided tests. **b**, Using out-of-sample likelihoods to do model comparison, the SR is the best-fitting model in terms of likelihood weights (similar to Akaike weights), the proportion of subjects best-fit and PXP. **c**, The simulations suggest that even a small amount of on-task SR replay helps an agent re-cache its representation of the social network after transition re-evaluation, resulting in improved navigation accuracy. The dashed horizontal lines reflect chance-level accuracy.

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[24,30,48], but that they also generate entirely new, synthetic 'walks' through the environment[29,49]. Moreover, the fact that we observe the greatest boost in navigational improvement for long-range problems is consistent with past findings demonstrating that sleep privileges memory abstraction[31–35]. We find that, after overnight rest, behaviour was consistent with larger SR gammas on the second day of testing, which aligns with prior work showing that sleep appears to be important for 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[15].

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 navigation performance improves overnight, but a fuller computational account of such a mechanism requires characterizing the content and amount of replay experienced. Past neurobiological findings strongly suggest that replay sequences consist of items that were experienced close together in time (for example, adjacent locations in a maze that are part of the same path)[24]. 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. It also remains unknown how much neural replay is necessary for improving decision making (in offline replay overnight and/or in on-task replay[47]), and whether a computational model could provide reasonable estimates of the amount of neural replay occurring in individual subjects. Finally, future research should test whether the benefits of SR replay are linked to sleep specifically or can also be observed with longer periods of awake rest.

A related question revolves around the neural instantiation of multi-step abstraction. Although we leverage the SR in this work, we note that multi-step abstraction is a much more general representational strategy that could be implemented using many mechanisms with varying degrees of biological plausibility[9,19,20]. Despite the SR depending heavily on the temporal dynamics of experience[14,50], multi-step abstraction appears to describe how people represent social networks even when observations of social interaction are temporally disjointed[4]. It is therefore possible that the SR successfully describes social network representation because it is a useful method for discovering structure[15], rather than being a faithful model of neural

 https://doi.org/10.1038/s41562-024-01990-w

computation. A particularly intriguing possibility is that the brain may encode components of a network's structure (that is, basis sets) that afford greater flexibility in assembling useful representations when navigating a variety of social environments[9]. There are many ways that the brain could perform inference over graphs[20,51–53], and it may be useful for future work to examine what kinds of basis sets are afforded by various methods of graph inference, including multi-step abstraction.

In summary, people can reason about information flow efficiently in social networks by caching knowledge about long-range connections in abstracted cognitive maps. Our results provide mechanistic insights into how these abstract cognitive maps are learned and shaped offline by a replay-like mechanism that allows the successful navigation over longer-range friendships.

Methods

Subjects

In study 1, we recruited $N = 50$ subjects (34 female, 15 male and 1 non-binary; mean age 20.6 years old, s.d. = 2.81 years). 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 1 was non-binary; the mean age was 23.1 years old, s.d. = 4.63 years. In study 3, we recruited $N = 50$ subjects, but lost four data points due to experimenter error, leaving a final sample size of $N = 46$ (30 female, 16 male; mean age 23.0 years old, s.d. = 4.46 years). All subjects received US $10 per hour as monetary compensation for their first study session. For the second study session, subjects in study 2 were paid US $15, and subjects in study 3 were paid US $20. Subjects in studies 2 and 3 could earn additional cash bonuses of up to US $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 (protocol 1607001555), and all subjects provided informed consent.

Overview

In study 1, a 1-day study (Fig. 1f), subjects first learned about a novel social network (Fig. 1a), completed a memory test (Fig. 1b) and then were tasked with solving social navigation problems (Fig. 1c) about the network they had just learned about (Fig. 1d). Details about each procedure are provided in subsequent sections.

Study 2 was a 2-day study (Fig. 1f), where day 1 was identical to study 1. On day 2, subjects returned to the laboratory 24 h after their first session. In this second session, subjects completed the same memory test and social navigation task as they had on day 1, then completed the social navigation task a third time after being informed about changes in network members' friendships (Fig. 1e).

Study 3 was a 2-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 min rest period at the end of day 1, after which subjects completed the memory test and social navigation task again (Fig. 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 (Fig. 1a). On each trial, subjects were shown one 'target' network member and were required to find all of the target's friends among the remaining 12 cards, which were initially displayed face down. The 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 (Fig. 1a). Once all of the target's friends were identified, subjects were given 3 s 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 min to complete. All face stimuli were drawn from the Chicago Face Database[54].

Memory test

Subjects completed a memory test immediately after the learning task (Fig. 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 min 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 (Fig. 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 min to complete on day 1 (Fig. 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 min and then completed the same navigation task again (Fig. 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 s. Across the entire 15 min 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 and 3 (Fig. 1f). Afterwards, we administered a transition re-evaluation procedure to test how alterations in network structure would impact navigation accuracy (Fig. 1f). Specifically, we instructed subjects that two individuals who had previously been friends were no longer friends (Fig. 1d,e) and that a new friendship had been formed between two other network members (Fig. 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.

Behavioural analysis

We used the R package glmmTMB to estimate mixed-effects logistic regression models[55]. Whenever appropriate, we pooled data across the three studies to maximize statistical power. To account for non-independent observations, the models included random intercepts for each subject, as well as a random intercept for each study. To test memory accuracy, we pooled data from studies 2 and 3 and estimated a model where study session (that is, day 1 versus day 2) was both a fixed-effects predictor and a random slope. In the social navigation task, shortest path distance was defined as the graph distance between the correct source and target after removing the 'message sender' from the network (that is, because the source was not permitted to 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 (that is, both answers were correct), 27 trials required online re-evaluation of path distance (that is, because the shortest possible path would have required sending the letter back to the sender) and 3 trials required both online re-evaluation and had the same shortest path distance. Our main analyses focussed on the remaining subset of 115 trials where there was an unambiguously correct answer.

We tested four behavioural hypotheses in total. Our first hypothesis tested whether subjects were able to solve navigation problems with above-chance accuracy on day 1, immediately after learning about the social network (Fig. 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 (that is, shortest paths of two, three and four), we iteratively re-parameterized the model by making each distance the reference category. Our second hypothesis tested whether subjects' navigation accuracy improved on day 2 after overnight rest (Fig. 1f), and therefore pooled data from studies 2 and 3. This model included fixed-effects predictors for shortest path distance, the study session (that is, day 1 versus 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. Our third hypothesis tested whether 15 min of awake rest would improve navigation accuracy. This model was functionally identical to the 2 day model, except that it used data from the study 3 navigation tasks before rest and after awake rest. Finally, our fourth hypothesis tested whether changes in the network (that is, transition re-evaluation) would result in attenuated navigation accuracy using a model that was similar to the 2 day model, but with two key differences: (1) the model compared navigation performance from before rest, after overnight rest and after transition re-evaluation and (2) to control for the possible confound that the transition re-evaluated trials were more difficult than the main set of navigation trials, the model included an additional predictor indexing the absolute difference in the two sources' shortest path distance to the target, which serves as a proxy for task difficulty (Supplementary Fig. 14).

SR

In its typical use in reinforcement learning, the 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 (that is, the lookahead $L = \frac{1}{1-\gamma}$). This has a straightforward translation to social navigation problems, such as 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 $M(s, t)$. Here, we use function notation to index row s and column t of the matrix M.

In all SR-replay simulations, our implementation used a standard delta-rule method to update M (equation (1) and Figs. 3c and 4c). 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 $M(s)$ (this notation refers 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 the friendships of the s and t 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 fixed to 0.1 following past work[4,16,23]. As friendships are bidirectional in our study, each learning event prompted two updates, one for s and another for t. In the simulations, we treated each novel observation as a single learning event. As the learning task consisted of six blocks, each containing learning events for both $s \to t$ and $t \to s$, the no-replay SR learned from 12 observations of each friendship.

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

In all parameter fitting, we used an analytic form of the SR to generate asymptotic representations (equation (2) and Fig. 2b–e), where I is the identity matrix, T is the transition matrix and X^{-1} refers to the matrix inverse.

$$M = \sum_{k=0}^{\infty} \gamma^k T^k = (I - \gamma T)^{-1}. \tag{2}$$

We modelled an agent's choice between two candidate sources using a softmax choice rule (equation (3)), 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 $\beta \to \infty$, and more stochastically as $\beta \to 0$.

$$P(\text{choose SourceA}) = \frac{\exp[\beta M(s,\ t\,|\,s = \text{SourceA})]}{\exp[\beta M(s,\ t\,|\,s = \text{SourceA})] + \exp[\ \beta M(s, t\,|\,s = \text{SourceB})}. \tag{3}$$

Model-based planning

To compare multi-step abstraction against model-based planning, we tested several psychologically plausible mechanisms through which an agent could perform online planning.

The BFS-forwards model assumes that people perform two forwards searches in parallel, one from each source, until the target appears in one of the searches. 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. It is therefore possible that the target may, by chance, first appear in the search associated with the incorrect answer, such that the asymptotic predictions of the BFS-forwards model fall short of perfect accuracy. The BFS-forwards model contains a single 'search threshold' parameter, which causes the agent to become increasingly likely to 'give up' and choose randomly when searches require iterating across long distances.

The BFS-backwards model assumes that people start a single search from the target, then iterate backwards through the network until one of the two sources is found. Backwards searches are normatively more efficient than forwards searches, and humans appear to prefer using backwards search when it is possible for them to do so[37]. If the backwards search is allowed to run to completion, this guarantees that the subject will choose the correct answer, as BFS always finds the shortest path distance between two nodes in a graph. Like BFS-forwards, our implementation of BFS-backwards contains a single search threshold parameter. Although the BFS-backwards and BFS-forwards models use the same underlying search mechanism, we note that they make very different predictions about human behaviour (Supplementary Fig. 1).

To our knowledge, there is no analytic method for computing choice and 'reaction time' distributions from a BFS-based search process, so we simulated how our BFS agent would solve each trial 5,000 times (description of algorithm in Supplementary Information). To fit parameters to subjects' behaviour, 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 (equation (4)). The search threshold parameter τ was estimated as the value at which a softmax with $\beta = 1$ was indifferent between choosing to complete a search (based on the average length of the BFS search for a given trial) and giving up (equation (5)). Likelihoods were weighted accordingly (equation (6)). 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.

$$P(\text{choose SourceA}|\text{completing BFS}) = \text{avg}(\text{choose SourceA}), \quad (4)$$

$$P(\text{completing BFS}) = \frac{\exp(\tau)}{\exp(\tau) + \exp(\text{avg}[\text{search length}])}, \quad (5)$$

$$\begin{aligned} &P(\text{choose SourceA}) \\ &= [P(\text{choose SourceA}|\text{completing BFS}) \times P(\text{completing BFS})] \\ &\quad +[0.5 \times (1 - P(\text{completing BFS}))]. \end{aligned} \quad (6)$$

Finally, the ideal observer model assumes that people are able to compute the shortest path distances in the graph, and subsequently chooses between the two options using a softmax choice rule. We note that there are a number of psychologically plausible processes an agent could use to compute shortest path distances, including backwards BFS. Our goal is not to adjudicate between different process models, but rather to test whether any such ideal observer could provide a compelling alternative explanation for our empirical results. This model contains a single parameter, the softmax inverse temperature β, which controls the agent's sensitivity to the sources' shortest path distances from the target (equation (7)).

$$\begin{aligned} &P(\text{choose SourceA}) \\ &= \frac{\exp[\beta\,\text{dist}(\text{SourceA}, \text{Target})]}{\exp[\beta\,\text{dist}(\text{SourceA}, \text{Target})] + \exp[\beta\,\text{dist}(\text{SourceB}, \text{Target})]}. \end{aligned} \quad (7)$$

Computational modelling

Maximum-likelihood parameter-fitting was performed using R's default optimizer, using the Nelder–Mead algorithm for the two-parameter SR and the Brent algorithm for all the single-parameter planning models. Parameters were fit independently for each subject. The SR model was re-estimated 25 times, keeping only the estimates that best maximized the likelihood. Single-parameter models were only estimated once, as the fitting procedure is akin to grid search. Plots of all estimated parameters can be found in Supplementary Fig. 4.

To verify that parameter estimates are psychologically meaningful, we performed parameter recovery and model confusability analyses for all models under consideration (Supplementary Fig. 5)[56]. Parameter recovery analyses indicate that all key parameters are straightforwardly interpretable, and the model confusability analyses confirm that the parameter fitting procedure is not biased in favour of our hypothesis.

In our primary model comparison procedure, we found converging evidence from three metrics: Akaike weights, the proportion of subjects best fit by a particular model and PXP. Two of these metrics are more descriptive: Akaike weights quantify the conditional probabilities for each model based on differences in AIC[39], and the proportion of subjects best-fit by a particular model provides an intuitive sense for models' goodness of fit. PXP provides a formal test of a model's group-level fit compared with other candidate models[38], and was computed using R software written by Matteo Lisi (https://github.com/mattelisi/bmsR). As our measure of log evidence, we used the Akaike information criterion corrected for relatively few data points (that is, AICc), which penalizes models in proportion to the number of free parameters estimated[39].

Using the parameters estimated from the primary trials of interest (detailed in 'Behavioural analysis'), we tested how well each model is able to explain held-out data by computing out-of-sample likelihoods. This was done by simulating the models' predictions given the estimated parameters, then computing the resulting logistic loss based on subjects' actual behaviours in the held-out data. After computing out-of-sample likelihoods for the transition re-evaluated navigation trials, we performed formal model comparison as detailed above, except using log-likelihoods rather than AICc, as no parameters were fit for out-of-sample prediction.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

All data needed to reproduce the analyses are available in a publicly accessible GitHub repository at https://github.com/feldmanhalllab/network-navigation-replay (ref. 57).

Code availability

All code needed to reproduce the analyses are available in a publicly accessible GitHub repository at https://github.com/feldmanhalllab/network-navigation-replay (ref. 57).

References

1. Travers, J. & Milgram, S. An experimental study of the small world problem. *Sociometry* **32**, 425–443 (1969).
2. Christakis, N. A. & Fowler, J. H. Social contagion theory: examining dynamic social networks and human behavior. *Stat. Med.* **32**, 556–577 (2013).
3. Son, J.-Y., Bhandari, A. & FeldmanHall, O. Cognitive maps of social features enable flexible inference in social networks. *Proc. Natl Acad. Sci. USA* **118**, e2021699118 (2021).
4. Son, J.-Y., Bhandari, A. & FeldmanHall, O. Abstract cognitive maps of social network structure aid adaptive inference. *Proc. Natl Acad. Sci. USA* **120**, e2310801120 (2023).
5. Tolman, E. C. Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189–208 (1948).
6. O'Keefe, J. & Nadel, L. *The Hippocampus as a Cognitive Map* (Clarendon, 1978).
7. 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).
8. Bellmund, J. L. S., Gärdenfors, P., Moser, E. I. & Doeller, C. F. Navigating cognition: spatial codes for human thinking. *Science* **362**, eaat6766 (2018).
9. Behrens, T. E. J. et al. What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron* **100**, 490–509 (2018).
10. Constantinescu, A. O., O'Reilly, J. X. & Behrens, T. E. J. Organizing conceptual knowledge in humans with a gridlike code. *Science* **352**, 1464–1468 (2016).
11. 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).
12. Tavares, R. M. et al. A map for social navigation in the human brain. *Neuron* **87**, 231–243 (2015).
13. 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).

14. Dayan, P. Improving generalization for temporal difference learning: the successor representation. *Neural Comput.* **5**, 613–624 (1993).
15. Momennejad, I. Learning structures: predictive representations, replay, and generalization. *Curr. Opin. Behav. Sci.* **32**, 155–166 (2020).
16. Momennejad, I. et al. The successor representation in human reinforcement learning. *Nat. Hum. Behav.* **1**, 680–692 (2017).
17. 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 Comput. Biol.* **13**, e1005768 (2017).
18. Stachenfeld, K. L., Botvinick, M. M. & Gershman, S. J. The hippocampus as a predictive map. *Nat. Neurosci.* **20**, 1643–1653 (2017).
19. Lynn, C. W. & Bassett, D. S. How humans learn and represent networks. *Proc. Natl Acad. Sci. USA.* **117**, 29407–29415 (2020).
20. Lynn, C. W., Kahn, A. E., Nyema, N. & Bassett, D. S. Abstract representations of events arise from mental errors in learning and memory. *Nat. Commun.* **11**, 2313 (2020).
21. Momennejad, I. & Howard, M. W. Predicting the future with multi-scale successor representations. Preprint at *bioRxiv* https://doi.org/10.1101/449470 (2018).
22. Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B. & Botvinick, M. M. Neural representations of events arise from temporal community structure. *Nat. Neurosci.* **16**, 486–492 (2013).
23. Pudhiyidath, A. et al. Representations of temporal community structure in hippocampus and precuneus predict inductive reasoning decisions. *J. Cogn. Neurosci.* **34**, 1736–1760 (2022).
24. Foster, D. J. Replay comes of age. *Annu. Rev. Neurosci.* **40**, 581–602 (2017).
25. 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. *Nat. Commun.* **9**, 3920 (2018).
26. Sun, W., Advani, M., Spruston, N., Saxe, A. & Fitzgerald, J. E. Organizing memories for generalization in complementary learning systems. *Nat. Neurosci.* **26**, 1438–1448 (2023).
27. 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).
28. Foster, D. J. & Wilson, M. A. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**, 680–683 (2006).
29. Igata, H., Ikegaya, Y. & Sasaki, T. Prioritized experience replays on a hippocampal predictive map for learning. *Proc. Natl Acad. Sci. USA* **118**, e2011266118 (2021).
30. Zhenglong, Z., Michael, J. K. & Anna, C. S. Replay as context-driven memory reactivation. Preprint at *bioRxiv* https://doi.org/10.1101/2023.03.22.533833 (2023).
31. Ellenbogen, J. M., Hu, P. T., Payne, J. D., Titone, D. & Walker, M. P. Human relational memory requires time and sleep. *Proc. Natl Acad. Sci. USA* **104**, 7723–7728 (2007).
32. Lewis, P. A. & Durrant, S. J. Overlapping memory replay during sleep builds cognitive schemata. *Trends Cogn. Sci.* **15**, 343–351 (2011).
33. Lutz, N. D., Diekelmann, S., Hinse-Stern, P., Born, J. & Rauss, K. Sleep supports the slow abstraction of gist from visual perceptual memories. *Sci. Rep.* **7**, 42950 (2017).
34. 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. *Sci. Rep.* **12**, 15086 (2022).
35. Klinzing, J. G., Niethard, N. & Born, J. Mechanisms of systems memory consolidation during sleep. *Nat. Neurosci.* **22**, 1598–1610 (2019).
36. 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 Comput. Biol.* **19**, e1011087 (2023).
37. Callaway, F. et al. Rational use of cognitive resources in human planning. *Nat. Hum. Behav.* **6**, 1112–1125 (2022).
38. Rigoux, L., Stephan, K. E., Friston, K. J. & Daunizeau, J. Bayesian model selection for group studies—revisited. *NeuroImage* **84**, 971–985 (2014).
39. Wagenmakers, E.-J. & Farrell, S. AIC model selection using Akaike weights. *Psychon. Bull. Rev.* **11**, 192–196 (2004).
40. Kurth-Nelson, Z., Economides, M., Dolan, R. J. & Dayan, P. Fast sequences of non-spatial state representations in humans. *Neuron* **91**, 194–204 (2016).
41. Schuck, N. W. & Niv, Y. Sequential replay of non-spatial task states in the human hippocampus. *Science* **364**, eaaw5181 (2019).
42. Gómez, R. L., Bootzin, R. R. & Nadel, L. Naps promote abstraction in language-learning infants. *Psychol. Sci.* **17**, 670–674 (2006).
43. Lau, H., Alger, S. E. & Fishbein, W. Relational memory: a daytime nap facilitates the abstraction of general concepts. *PLoS ONE* **6**, e27139 (2011).
44. Pereira, S. I. R. et al. Rule abstraction is facilitated by auditory cuing in REM sleep. *J. Neurosci.* **43**, 3838–3848 (2023).
45. St Clair, M. C. & Monaghan, P. Language abstraction: consolidation of language structure during sleep. In *Proc. of the Annual Meeting of the Cognitive Science Society* **30** (Cognitive Science Society, 2008).
46. Walker, M. P. & Stickgold, R. Overnight alchemy: sleep-dependent memory evolution. *Nat. Rev. Neurosci.* **11**, 218 (2010).
47. Wittkuhn, L., Krippner, L. M. & Schuck, N. W. Replay in human visual cortex is linked to the formation of successor representations and independent of consciousness. Preprint at *bioRxiv* https://doi.org/10.1101/2022.02.02.478787 (2022) .
48. Jadhav, S. P., Kemere, C., German, P. W. & Frank, L. M. Awake hippocampal sharp-wave ripples support spatial memory. *Science* **336**, 1454–1458 (2012).
49. Stoianov, I., Maisto, D. & Pezzulo, G. The hippocampal formation as a hierarchical generative model supporting generative replay and continual learning. *Prog. Neurobiol.* **217**, 102329 (2022).
50. Gershman, S. J., Moore, C. D., Todd, M. T., Norman, K. A. & Sederberg, P. B. The successor representation and temporal context. *Neural Comput.* **24**, 1553–1568 (2012).
51. Lau, T., Gershman, S. J. & Cikara, M. Social structure learning in human anterior insula. *Elife* **9**, e53162 (2020).
52. Whittington, J. C. R. et al. The Tolman–Eichenbaum machine: unifying space and relational memory through generalization in the hippocampal formation. *Cell* **183**, 1249–1263.e1223 (2020).
53. Wu, C. M., Schulz, E. & Gershman, S. J. Inference and search on graph-structured spaces. *Comput. Brain Behav.* **4**, 125–147 (2021).
54. Ma, D. S., Correll, J. & Wittenbrink, B. The Chicago face database: a free stimulus set of faces and norming data. *Behav. Res. Methods* **47**, 1122–1135 (2015).
55. Brooks, M. E. et al. glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling. *R J.* **9**, 378–400 (2017).
56. Wilson, R. C. & Collins, A. G. E. Ten simple rules for the computational modeling of behavioral data. *Elife* **8**, e49547 (2019).
57. Network-navigation-replay. *GitHub* https://github.com/feldmanhalllab/network-navigation-replay (2024).

Acknowledgements

We thank I. Aslarus, K. Danowski, E. Duchan, Y.-F. Jerry Hu, A. Lawrence, J. Palfy, V. Poyraz, M. Malhotra, M. Mazumder, S. Shulman, A. Stein, S. Vaca Narvaja and J. Wang for assisting with data collection. We thank A. Maddah for developing some of the task code used in these studies and Y. Yang Teoh for helpful computational modelling advice. Part of this research was conducted using computational resources and services at the Centre 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.). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the paper.

Author contributions

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

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41562-024-01990-w.

Correspondence and requests for materials should be addressed to Apoorva Bhandari or Oriel FeldmanHall.

Peer review information *Nature Human Behaviour* thanks Charley Wu and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

nature portfolio

Corresponding author(s): Oriel FeldmanHall, Apoorva Bhandari
Last updated by author(s): Aug 12, 2024

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☒	☐	Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Data were collected using code written in the jsPsych 6.2.0 library.
Data analysis	Data were analyzed using R 4.3.1. All analysis code and documentation is available in a publicly-accessible GitHub repository: https://github.com/feldmanhalllab/network-navigation-replay

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

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

Experimental stimuli included photographs of human faces drawn from the Chicago Face Database: https://www.chicagofaces.org/

Research involving human participants, their data, or biological material

Policy information about studies with human participants or human data. See also policy information about sex, gender (identity/presentation), and sexual orientation and race, ethnicity and racism.

Reporting on sex and gender	Our demographic information includes participants' self-reported gender identity. No analyses were performed using gender variables, as the phenomena of interest are not expected to differ based on sex or gender identity.
Reporting on race, ethnicity, or other socially relevant groupings	No analyses were performed using race, ethnicity, or other socially relevant groupings.
Population characteristics	See details in "Behavioural & social sciences study design" section.
Recruitment	Participants were recruited using non-targeted ads distributed among the university and the surrounding community. This recruitment strategy, common in psychology research, is likely to draw heavily from a 'WEIRD' population (Western, Educated, Industrialized, Rich, and Democratic), which is known not to be representative of many human populations. 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.
Ethics oversight	The study protocol was approved by Brown University's Institutional Review Board (Protocol 1607001555), and all subjects provided informed consent.

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☒ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	We used a quantitative experimental study design.
Research sample	The research sample consisted of undergraduate and graduate students at Brown University, as well as adult volunteers from the surrounding community. 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). The research sample reflects a convenience sampling strategy, and therefore may not be a representative sample.
Sampling strategy	We used a convenience sample in this work. Using G*Power, we determined that we would need a sample of N = 41 to achieve 80% power using a paired t-test with an assumed effect size of Cohen's d = 0.4. We therefore aimed to recruit N = 50 for each of the three studies. We further maximized statistical power by pooling data from the three studies in a mixed-effects regression model whenever possible.
Data collection	We collected data using a computerized task. Participants were allowed to complete the task in a private room while the researchers waited in an outside room in case participants had any questions or concerns. Researchers were not blind to the experimental condition or study hypothesis during data collection.
Timing	Data for study 1 were collected from November 23, 2021 through March 25, 2022. Data for study 2 were collected from April 21, 2022 through June 17, 2022. Data for study 3 were collected from February 22, 2023 through June 9, 2023.
Data exclusions	No data were excluded from the analysis, though four datapoints were lost in study 3 due to experimenter error (see above).
Non-participation	No participants dropped out or declined participation.
Randomization	We used a within-subjects study design, and therefore no randomization was used.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☒	☐ Animals and other organisms
☒	☐ Clinical data
☒	☐ Dual use research of concern
☒	☐ Plants

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Plants

Seed stocks	*Report on the source of all seed stocks or other plant material used. If applicable, state the seed stock centre and catalogue number. If plant specimens were collected from the field, describe the collection location, date and sampling procedures.*
Novel plant genotypes	*Describe the methods by which all novel plant genotypes were produced. This includes those generated by transgenic approaches, gene editing, chemical/radiation-based mutagenesis and hybridization. For transgenic lines, describe the transformation method, the number of independent lines analyzed and the generation upon which experiments were performed. For gene-edited lines, describe the editor used, the endogenous sequence targeted for editing, the targeting guide RNA sequence (if applicable) and how the editor was applied.*
Authentication	*Describe any authentication procedures for each seed stock used or novel genotype generated. Describe any experiments used to assess the effect of a mutation and, where applicable, how potential secondary effects (e.g. second site T-DNA insertions, mosiacism, off-target gene editing) were examined.*