

Novel Inferences in a Multidimensional Social Network Use a Grid-like Code

Seongmin A. Park¹, Douglas S. Miller^{1,2}, Erie D. Boorman^{1,3}

¹. Center for Mind and Brain, University of California, Davis, USA

². Center for Neuroscience, University of California, Davis, USA

³. Department of Psychology, University of California, Davis, USA

Novel Inferences in a Multidimensional Social Network Use a Grid-like Code

ABSTRACT

Generalizing experiences to guide decision making in novel situations is a hallmark of flexible behavior. It has been hypothesized such flexibility depends on a cognitive map of an environment or task, but directly linking the two has proven elusive. Here, we find that discretely sampled abstract relationships between entities in an unseen two-dimensional (2-D) social hierarchy are reconstructed into a unitary 2-D cognitive map in the hippocampus and entorhinal cortex. We further show that humans utilize a grid-like code in several brain regions, including entorhinal cortex and medial prefrontal cortex, for *inferred* direct trajectories between entities in the reconstructed abstract space during discrete decisions. Moreover, these neural grid-like codes in the entorhinal cortex predict neural decision value computations in the medial prefrontal cortex and temporoparietal junction area during choice. Collectively, these findings show that grid-like codes are used by the human brain to infer novel solutions, even in abstract and discrete problems, and suggest a general mechanism underpinning flexible decision making and generalization.

Humans and other animals have a remarkable capacity for inferring novel solutions in situations never encountered before. This extreme flexibility sets humans in particular apart from artificial intelligence, where inferring new solutions from sparse data and flexibly generalizing knowledge across tasks remains a significant challenge (Timothy E.J. Behrens et al., 2018; Sutton & Barto, 1998). As early as the 1940s, it was originally proposed that animals form and use some type of a “cognitive map” of the environment that allows them to flexibly infer novel routes and shortcuts through a spatial environment to reach a known reward’s location (Tolman, 1948).

In the context of spatial navigation, so-called “place cells” in the hippocampus (HC) and “grid cells” in medial entorhinal cortex (EC) have been proposed to represent the animal’s current or predicted position in a spatiotemporal environment (Ekstrom & Ranganath, 2018; Hafting, Fyhn, Molden, Moser, & Moser, 2005; O’Keefe & Nadel, 1978; Stachenfeld, Botvinick, & Gershman, 2017) and to provide a spatial metric of that environment (Bush, Barry, Manson, & Burgess, 2015), respectively. Together with related cell types (e.g. border cells, band cells, etc.), these cells may provide a neuronal basis for a cognitive map of space. Critically, recent biologically inspired computational models have further proposed how these grid codes in EC may reflect a basis for statistical transitions in 2-D topologies that can in principle enable animals and artificial agents to compose new routes and find shortcuts to reach goals in spatial environments (Banino et al., 2018;

Timothy E.J. Behrens et al., 2018; Bush et al., 2015; Whittington et al., 2019). These findings suggest a novel framework for understanding how humans might effectively “shortcut” solutions in new situations.

Representing non-spatial relationships between entities or task states in the world as a cognitive map would likewise be incredibly powerful because it could theoretically be leveraged to make inferences from sparse experiences that can dramatically accelerate learning and even guide decisions never faced before (Timothy E.J. Behrens et al., 2018; Kriete, Noelle, Cohen, & O'Reilly, 2013; Wang et al., 2018). This is because, in theory, a cognitive map of abstract relationships, such as a conceptual space (Constantinescu, O'Reilly, & Behrens, 2016) or task space (Schuck, Cai, Wilson, & Niv, 2016; Zhou et al., 2019), would allow for “novel routes” and “shortcuts” to be inferred, as in physical space (Timothy E.J. Behrens et al., 2018; Whittington et al., 2019) (**Fig. 1A**). Notably, such direct novel inferences go beyond chaining together directly experienced associations (**Fig. 1A**), and are seen as the defining characteristic of a genuine cognitive map of spatial environments (Bennett, 1996; O'Keefe & Nadel, 1978; Tolman, 1948). Despite these theoretical insights, whether such direct inferences (as opposed to chaining together previously experienced associations) are in fact used for non-spatial decision making, and whether a grid code underpins these inferences about non-spatial relationships, is unknown.

Recent pioneering studies have identified hippocampal place-like and entorhinal grid-like neural modulation for behaviorally relevant, continuous, non-spatial stimulus dimensions, including sound frequency (Aronov, Nevers, & Tank, 2017), metric length of bird legs and necks in a conceptual “bird space”(Constantinescu et al., 2016), and odor concentration (Bao et al., 2019). These findings support theories proposing that place and grid cells revealed in physical space may reflect a general mechanism for organizing relationships between any behaviorally relevant dimensions (Timothy E.J. Behrens et al., 2018; Eichenbaum & Cohen, 2014). However, these recent studies have all used *continuous* stimulus dimensions presented to the subject during task performance (e.g. visual, auditory, olfactory), thereby providing continuous feedback about one's “position” along the directly experienced trajectories (Bao et al., 2019; Constantinescu et al., 2016; Doeller, Barry, & Burgess, 2010; Nau, Navarro Schröder, Bellmund, & Doeller, 2018; Stachenfeld et al., 2017). Because the continuous trajectories in these studies were experienced during training or the task, behavior did not require novel inferences. In the real world, a wide range of everyday decisions involve discrete entities that vary along multiple abstract dimensions whose feature values must be learned discretely from piecemeal experiences, and in the absence of explicit sensory feedback. Here, we test the hypotheses that, first, such piecemeal experiences are integrated into an abstract cognitive map in the HC–EC system, and,

second, a grid code is used to make direct inferences to guide decision making based on that cognitive map (see **Fig. 1B,C**).

RESULTS

We leveraged established fMRI approaches (Constantinescu et al., 2016; Doeller et al., 2010; Kriegeskorte, Mur, & Bandettini, 2008) to test for a reconstructed cognitive map of an abstract two-dimensional (2-D) social hierarchy, and a grid-like code for inferred direct trajectories between entities during decision making (**Fig.1B**). Importantly, the social hierarchy was never seen but could be reconstructed from the outcomes of binary comparisons about the higher rank individual between pairs on one dimension at a time, with each dimension learned on a separate day. To test our hypotheses, we designed a “partner selection” task that required participants (n=21) to infer a better collaborator for a given “entrepreneur” by comparing their relative ordinal ranks in two ecologically important social dimensions: their competence and popularity (Fiske, Cuddy, & Glick, 2007). During behavioral training, participants learned the relative rank between individuals in the two independent dimensions through feedback-based binary comparisons. Participants only learned about individuals who differed by one rank level and on one dimension at a time (**Fig. 1D**), with each dimension learned on a separate day (on average 2 days apart). Importantly, participants were never shown the true hierarchy, nor were they asked to construct it spatially, but they could theoretically reconstruct it through transitive inferences within a dimension on a given day (Kumaran, Melo, & Duzel, 2012) and integration across dimensions learned on different days (**Fig. S1**). During fMRI participants were asked to choose a better partner for a given “entrepreneur” (F0) between two potential partners (F1 or F2) (**Fig. 1E**). Critically, these individuals’ 2-D rank in the combined social space was never learned during training, meaning that each decision during the fMRI task required novel inferences to be made by combining the two separately learned dimensions.

In our task the better collaboration was determined by the growth potential (GP) of each pair. To compute the GP, participants were instructed they should weight the ranks of the pair in both dimensions equally, with the rank of a pair in each dimension determined by the higher rank individual between the two. Therefore, the GP corresponds to the area drawn by the higher rank of the two entrepreneurs in each dimension in the 2-D hierarchy (**Fig. 1F**). We hypothesized people would infer direct trajectories over the reconstructed 2-D space between the positions of the entrepreneur, F0, and each potential partner, F1 and F2, to compute the GP for each collaboration (**Fig. 1B,E,F**). After training on the partner selection task, performed with a different choice set and in the absence of any choice feedback, participants who passed our performance criterion (>95% accuracy, see Methods)

were highly accurate in choosing the better collaboration during the fMRI session (mean \pm SE accuracy = $99.04 \pm 0.21\%$).

We first tested whether the HC-EC system would reconstruct the true unseen social hierarchy as a 2-D cognitive map. Specifically, we reasoned that more proximal people in the true 2-D social space (measured by their Euclidian distance) would be represented progressively and linearly more similarly (**Fig. 2A**). Using representational similarity analysis (RSA) applied to BOLD activity patterns evoked by each face presentation (at F1 and F2), we found that the level of pattern dissimilarity across voxels (measured by the Mahalanobis distance) in *a priori* anatomically defined regions of interests (ROIs) in the HC (Yushkevich et al., 2015) and EC (Amunts et al., 2005; Zilles & Amunts, 2010) (**Fig. 2B**) was robustly related to the pairwise Euclidian distance between individuals in the true 2-D social space ($p < 0.001$ FWE-corrected over ROIs) (**Fig. 2C**), such that closer people in the true 2-D space were represented increasingly more similarly. To ensure these pattern similarity effects in the HC and EC were better accounted for by the 2-D Euclidian distance than either 1-D distance alone, we decomposed the Euclidian distance into separate popularity and competence dimensions. We found significant effects of the 1-D rank distance in both popularity and competence dimensions on pattern similarity in HC and EC (popularity: $\tau_A = 0.09$; competence: $\tau_A = 0.08$; both $p < 0.001$ FWE-corrected; see **Fig. 2D** for alternative definitions of X and Y dimensions), consistent with a 2-D cognitive map. To further probe the representational architecture in HC and EC in a 2-D space, we applied multidimensional scaling on their activity patterns, which revealed a consistent relationship between the true Euclidian distances and the estimated distances between individuals' locations in a 2-D space (**Fig. 2E**; **Fig. S1E**). These analyses show that the HC-EC system integrates piecemeal, discrete observations about abstract relationships, learned on separate days from the outcomes of choices, into a unitary 2-D cognitive map.

Having found evidence that the 2-D social hierarchy is represented as a 2-D cognitive map in HC and EC, we next searched for established fMRI markers of a grid-like code, but for inferred, rather than previously experienced, trajectories during decision making (Constantinescu et al., 2016). Specifically, we searched for neural evidence for hexadirectional modulation for inferred trajectories through the reconstructed cognitive map (**Fig. 1B**) that we hypothesized were used to compute the growth potential to guide decisions: $\overrightarrow{F0F1}$ at the time of F1 presentation and $\overrightarrow{F0F2}$ at the time of F2 presentation (**Fig. 1E,F**). To test this hypothesis, we employed a two-step approach: first, we functionally identified ROIs showing hexagonal modulation for inferred trajectories independently of a participant's EC grid orientation for future independent tests, and, second, we search across

the whole brain for regions showing hexagonal modulation that is aligned to a participant's EC grid orientation.

For the first step, we tested for regions where the BOLD signal was hexagonally modulated (i.e. any linear combination of $\sin(6\theta)$ and $\cos(6\theta)$ using a Z-transformed F-statistic (as in (Constantinescu et al., 2016)), where θ is the inferred trajectory angle between two faces in the social space). Note that this analysis was performed to functionally identify independent EC ROIs from which to compute the grid orientation angle ϕ for a given subject for future tests. We found evidence of hexagonal modulation, independently of grid orientation, in the EC (at $Z > 2.3$, $p < 0.01$) and additionally in mPFC, posterior cingulate cortex (PCC), and posterior parietal cortex (PPC) ($p < 0.05$ whole-brain cluster corrected using permutation-based threshold-free cluster enhancement (TFCE)) (Smith & Nichols, 2009) (**Fig. S2A,B; Table S1**).

We then tested for hexagonal modulation aligned to the EC grid orientation for each participant using an unbiased cross-validation (CV) procedure from fMRI sessions acquired in the same day. To obviate concerns of selection bias, we separated a 'training dataset' for estimating the putative grid angle (ϕ) from the EC ROI functionally identified by step 1 above from a 'testing dataset' on which the effects of hexagonal symmetry were tested (using the regressor $\cos(6[\theta_t - \phi])$) (see Methods and **Fig. S2C**). For this critical test, we observed hexagonal modulation for inferred trajectories aligned to the estimated grid orientation in bilateral EC and mPFC ($p < 0.05$ TFCE-corrected; **Fig. 3; Table S2**; see **Fig. S3A** for confirmatory analyses on independent anatomically defined EC (Amunts et al., 2005; Zilles & Amunts, 2010) and mPFC (Neubert, Mars, Sallet, & Rushworth, 2015) ROIs). Importantly, control analyses indicated that the modulation in EC and mPFC was specific to a six-fold periodicity, as significant modulation was not seen for four-, five-, seven-, or eightfold periodicities (all $p > 0.05$). Furthermore, the effect at six-fold was significantly stronger than the control periodicities (all $p < 0.05$) (**Fig. 3B**).

Outside of these regions predicted from prior studies (Constantinescu et al., 2016; Doeller et al., 2010), we also observed within-day hexagonal modulation aligned to the EC grid orientation in bilateral temporo-parietal junction (TPJ) areas and bilaterally along the superior temporal sulci (STS) (**Fig. 3A**; see **Fig. S3B** for additional analyses of fusiform face area (FFA) ROIs). Notably, the TPJ, STS, and dorsal aspect of mPFC, in which we observed hexagonal modulation, are commonly recruited during tasks engaging theory of mind in social cognition (T. E. J. Behrens, Hunt, & Rushworth, 2009; Frith & Frith, 2012; Hampton, Bossaerts, & O'Doherty, 2008; Platt, Seyfarth, & Cheney, 2016; Wittmann, Lockwood, & Rushworth, 2018), including the modeling or simulation of different perspectives (Nicolle et al., 2012). Since participants were asked to take the perspective of the entrepreneur (F0)

and decide with whom they should partner, perspective-taking on the cognitive map was a likely requirement to perform the partner selection task successfully. This requirement may therefore explain their recruitment, and alignment with the grid code, in our partner selection task.

To test whether grid alignment was consistent for the same behavioral task over time, we acquired another fMRI dataset on the same subjects ($n=21$) on the same task more than a week later (average duration = 8.95 ± 0.81 days). We found hexagonal modulation aligned to the EC grid calculated from the other (left-out) day in a similar set of brain regions including the EC and mPFC, and also in PCC ($p < 0.05$ TFCE-corrected) (**Fig. 3C**; **Fig. S2D**; **Table S2**). Once again, the effect at the six-fold periodicity was significantly stronger than all other control periodicities (all $p < 0.05$). This set of brain regions – EC, mPFC, and PCC – that showed a consistent grid-like code over time for novel inferences through a reconstructed social space during decision making align with previous findings for trajectories actually traveled through physical space to remembered locations in virtual reality (Doeller et al., 2010).

To make optimal decisions in the partner selection task, participants need to compute and compare the growth potentials (GP) of each pair. The value of each decision option therefore corresponds to each pair's GP. A whole-brain analysis of parametric effects of the GP at the time of face 1 and face 2 presentations revealed significant effects of the GP in bilateral TPJ, mPFC, and EC ($p < 0.05$ TFCE-corrected; **Fig. 4A**; **Table S3**), indicating that BOLD activity in these regions quantitatively reflected the GP of a pair under consideration. Importantly, the GP and hexagonal modulation ($\cos(6[\theta_t - \phi])$) regressors showed little correlation (mean $r \pm SE = 0.04 \pm 0.03$), and both effects remained after including the other term as a regressor in the same GLM, indicating both terms explained independent variance in these regions (**Fig. S4**).

We hypothesized that the grid code for inferred vectors may be used in the computation of the GP based on a cognitive map in our task. If true, we predicted that the GP effects in TPJ and mPFC regions would be modulated by trajectories that were aligned compared to misaligned to the EC grid orientation. Examining the TPJ and mPFC ROIs functionally defined using the statistically independent GP regressor (at $p < 0.001$ uncorrected) showed greater GP effects for trajectories that were aligned compared to misaligned *with the EC grid orientation* in these regions (**Fig. 4B**). Notably, while there was a significant GP effect in TPJ and mPFC for both trajectories that were aligned and misaligned to the grid orientation, the effect was modulated such that it was significantly stronger for aligned trajectories (**Fig. 4D,E**; **Table S3**), consistent with a dependence of GP coding on EC grid

coding. Control analyses confirmed specificity for this 6-fold periodicity in TPJ and mPFC over control periodicities (all $p < 0.05$) (**Fig. 4B**).

We further reasoned that if the grid code for inferred trajectories is related to the process of constructing the values of decision options (here, the pair's GP), then those subjects with greater hexagonal modulation in EC would also be predicted to show greater neural GP effects. Consistent with this prediction, we found a positive correlation between the EC grid modulation effect and the GP effect in TPJ (**Fig. 4C**; $r = 0.58$, $p = 0.006$ for left TPJ and $r = 0.44$, $p = 0.05$ for right TPJ), but not in mPFC ($p > 0.05$), across participants, after controlling for choice accuracy.

At the time the second face (F2) is presented, subjects have all the information needed to make a decision based on a comparison of the two potential collaborations' GPs. Considerable evidence has supported a role for the ventromedial prefrontal cortex (vmPFC) in representing and comparing options' reward values for goal-directed decisions, a signal that also reflects decision confidence (T. E. J. Behrens, Hunt, Woolrich, & Rushworth, 2008; Boorman, Behrens, Woolrich, & Rushworth, 2009; Boorman, Behrens, & Rushworth, 2011; Clithero & Rangel, 2013; Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; De Martino, Fleming, Garrett, & Dolan, 2013; FitzGerald, Seymour, & Dolan, 2009; Hunt et al., 2012; Kolling, Behrens, Mars, & Rushworth, 2012; Lebreton, Abitbol, Daunizeau, & Pessiglione, 2015; Lim, O'Doherty, & Rangel, 2011; Noonan, Mars, & Rushworth, 2011; Park, Sestito, Boorman, & Dreher, 2019; Yamada, Louie, Tymula, & Glimcher, 2018). Searching for the equivalent value comparison term in our task— a comparison between the two options' values, or $|GP1 - GP2|$ — replicated this effect in vmPFC, and also revealed a significant effect in the EC extending into the HC ($p < 0.05$ TFCE-corrected) (**Fig. 4F**; **Table S3**). This effect in EC/HC is consistent with the interpretation that the comparison between options' values during decision making additionally involves the interconnected HC–EC system (Witter, Wouterlood, Naber, & Van Haeften, 2000) when it is based on a cognitive map (Bakkour et al., 2019).

DISCUSSION

We have shown that abstract and discrete relational structures, never seen but inferred from discrete samples through a series of binary comparisons on each dimension made on different days, are reconstructed into a 2-D cognitive map in the HP-EC system. Specifically, we found that the closer faces were ranked in the true 2-D social hierarchy, the increasingly more similar their neural activity patterns were in both HC and EC. This finding suggests that each entity is located according to its feature values along relevant dimensions in a 2-D relational space (Timothy E.J. Behrens et al., 2018; Eichenbaum & Cohen, 2014). Furthermore, we found evidence for hexagonal modulation for inferred direct

trajectories over the abstract social space during discrete decision making, further supporting the interpretation that people infer over a 2-D cognitive map of abstract relationships. Our findings thus suggest that having an explicit representation of structural relationships is beneficial for generalizing experiences that can be used to guide novel inferences during decision making.

We found evidence for a grid-like code in EC, mPFC, PCC, TPJ, and STS, for inferred direct trajectories through the reconstructed 2-D social space during decision making. This finding provides novel evidence that humans can use direct inferred trajectories in non-spatial decision making, a defining feature of a cognitive map of space (Bennett, 1996; O'Keefe & Nadel, 1978). Taken together with recent findings in non-social tasks (Bao et al., 2019; Constantinescu et al., 2016; Doeller et al., 2010; Jacobs et al., 2013; Nau et al., 2018), these findings suggest that grid-like codes are present in a markedly common set of core cortical brain regions including the EC and interconnected mPFC and PCC (Eichenbaum, 2017; Preston & Eichenbaum, 2013; Witter et al., 2000) for both spatial and non-spatial relational tasks. Furthermore, the grid-like code's alignment between EC and mPFC was consistent when performing the same task over a week later, consistent with a previous study (Constantinescu et al., 2016), indicating a relatively stable representation over time, even in the absence of practice. Here, we extend these previous reports by showing for the first time that first, a grid-like code can be used for discretely sampled abstract spaces and second, to compose novel inferences during non-spatial decision making.

On the other hand, our findings also suggest that additional brain regions – TPJ, dorsal mPFC, and STS (and also to some extent FFA (see Fig. S3B)) – may come online in shared alignment with the EC grid code, depending on the specific processes required by a particular task. Notably, these areas are thought to play an important role in theory of mind (T. E J Behrens et al., 2009; Frith & Frith, 2012; Hampton et al., 2008; Platt et al., 2016; Wittmann et al., 2018), which frequently involves simulating another person's mental state (Nicolle et al., 2012). Since participants were explicitly asked to select a partner on behalf of another entrepreneur, a manipulation that served to anchor trajectories to different locations in the social space, perspective-taking on the cognitive map was engaged in our task. Whether the grid-like modulation of the BOLD signal in these regions reflects grid cells locally outside of EC (Jacobs et al., 2013; Long & Zhang, 2018), or synchronized input between EC and other interconnected cortical areas (Eichenbaum, 2017; Preston & Eichenbaum, 2013; Witter et al., 2000), remains an intriguing open question for future studies.

Finally, our study provides novel evidence linking a grid-like code to the construction of decision values in mPFC and TPJ during decision making. Specifically, we found that the

GP of a pair, equivalent to the decision option's value in our task, was encoded in the mPFC and TPJ, after controlling for the hexagonal modulation effect. This finding is consistent with past studies showing these areas encode another person's decision value when predicting their decision (Nicolle et al., 2012; Piva et al., 2019). Importantly, this effect was modulated by the inferred trajectory's alignment to the EC grid orientation such that the value was encoded more strongly for aligned than misaligned trajectories. Notably, hexagonal modulation of theta band local field data from pre-surgical recordings has been identified for traveled trajectories in virtual reality (Maidenbaum, Miller, Stein, & Jacobs, 2018; Staudigl et al., 2018). Moreover, recent evidence indicates a key role for theta-phase synchrony between the hippocampal formation and orbitofrontal cortex that drives stimulus value coding in orbitofrontal cortex during value-based decision making (Knudsen & Wallis, 2020). Thus, one interpretation of the dependence of value coding in mPFC and TPJ on EC grid alignment in our study is that the grid code is theta-phase locked to the coding of the growth potential in mPFC and TPJ, thus providing an information channel between EC and these regions. Moreover, we found that those individuals who showed a stronger hexagonal modulation in EC also showed stronger GP effects in TPJ, even after controlling for choice accuracy. Collectively, these findings suggest that the grid code plays a central role in value construction when it is based on inferences over a cognitive map.

Taken together, our findings suggest the cognitive map in HC and EC provides a systematic framework which the brain can utilize for vector navigation through spaces defined by abstract relationships, allowing discrete entities or items to be related to each other in new situations, and thus generalization and the discovery of new solutions for choice. We suggest the alignment of the posteromedial network, and potentially other regions, with the entorhinal grid system may reflect a general mechanism for knowledge-guided inference during learning and decision making.

REFERENCES

- Amunts, K., Kedo, O., Kindler, M., Pieperhoff, P., Mohlberg, H., Shah, N. J., ... Zilles, K. (2005). Cytoarchitectonic mapping of the human amygdala, hippocampal region and entorhinal cortex: Intersubject variability and probability maps. *Anatomy and Embryology*, 210(5–6), 343–352. <https://doi.org/10.1007/s00429-005-0025-5>
- Aronov, D., Nevers, R., & Tank, D. W. (2017). Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature*, 543(7647), 719–722. <https://doi.org/10.1038/nature21692>
- Bakkour, A., Palombo, D. J., Zylberberg, A., Kang, Y. H. R., Reid, A., Verfaellie, M., ... Shohamy, D. (2019). The hippocampus supports deliberation during value-based decisions. *ELife*, 8. <https://doi.org/10.7554/eLife.46080>
- Banino, A., Barry, C., Uria, B., Blundell, C., Lillicrap, T., Mirowski, P., ... Kumaran, D. (2018). Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705), 429–433. <https://doi.org/10.1038/s41586-018-0102-6>
- Bao, X., Gjorgieva, E., Shanahan, L. K., Howard, J. D., Kahnt, T., & Gottfried, J. A. (2019). Grid-like Neural Representations Support Olfactory Navigation of a Two-Dimensional Odor Space. *Neuron*, 102(5), 1066–1075.e5. <https://doi.org/10.1016/j.neuron.2019.03.034>
- Behrens, T. E. J. J., Hunt, L. T., Woolrich, M. W., & Rushworth, M. F. S. S. (2008). Associative learning of social value. *Nature*, 456(7219), 245–249. <https://doi.org/10.1038/nature07538>
- Behrens, T. E. J., Hunt, L. T., & Rushworth, M. F. S. (2009). The computation of social behavior. *Science*, 324(5931), 1160–1164. <https://doi.org/10.1126/science.1169694>
- Behrens, Timothy E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron*, 100(2), 490–509. <https://doi.org/10.1016/j.neuron.2018.10.002>
- Bennett, A. T. D. (1996). Do animals have cognitive maps? *Journal of Experimental Biology*, 199(1), 219–224.
- Berens, P. (2009). CircStat: A MATLAB Toolbox for Circular Statistics. *Journal of Statistical Software*, 31(10). <https://doi.org/10.18637/jss.v031.i10>
- Boorman, E. D., Behrens, T. E. J., Woolrich, M. W., & Rushworth, M. F. S. (2009). How Green Is the Grass on the Other Side? Frontopolar Cortex and the Evidence in Favor of Alternative Courses of Action. *Neuron*, 62(5), 733–743. <https://doi.org/10.1016/j.neuron.2009.05.014>
- Boorman, E. D., Behrens, T. E., & Rushworth, M. F. (2011). Counterfactual choice and learning in a Neural Network centered on human lateral frontopolar cortex. *PLoS Biology*, 9(6), e1001093. <https://doi.org/10.1371/journal.pbio.1001093>
- Bush, D., Barry, C., Manson, D., & Burgess, N. (2015). Using Grid Cells for Navigation. *Neuron*, 87(3), 507–520. <https://doi.org/10.1016/j.neuron.2015.07.006>
- Clithero, J. A., & Rangel, A. (2013). Informatic parcellation of the network involved in the computation of subjective value. *Social Cognitive and Affective Neuroscience*. <https://doi.org/10.1093/scan/nst106>
- Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. J. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, 352(6292), 1464–1468. <https://doi.org/10.1126/science.aaf0941>
- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095), 876–879. <https://doi.org/10.1038/nature04766>
- De Martino, B., Fleming, S. M., Garrett, N., & Dolan, R. J. (2013). Confidence in value-based choice. *Nature Neuroscience*, 16(1), 105–110. <https://doi.org/10.1038/nn.3279>
- Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, 463(7281), 657–661. <https://doi.org/10.1038/nature08704>
- Eichenbaum, H. (2017). Prefrontal-hippocampal interactions in episodic memory. *Nature*

- Reviews Neuroscience*, 18(9), 547–558. <https://doi.org/10.1038/nrn.2017.74>
- Eichenbaum, H., & Cohen, N. J. (2014). Can We Reconcile the Declarative Memory and Spatial Navigation Views on Hippocampal Function? *Neuron*, 83(4), 764–770. <https://doi.org/10.1016/j.neuron.2014.07.032>
- Ekstrom, A. D., & Ranganath, C. (2018). Space, time, and episodic memory: The hippocampus is all over the cognitive map. *Hippocampus*, 28(9), 680–687. <https://doi.org/10.1002/hipo.22750>
- Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2007). Universal dimensions of social cognition: warmth and competence. *Trends in Cognitive Sciences*, 11(2), 77–83. <https://doi.org/10.1016/j.tics.2006.11.005>
- FitzGerald, T. H. B., Seymour, B., & Dolan, R. J. (2009). The role of human orbitofrontal cortex in value comparison for incommensurable objects. *Journal of Neuroscience*, 29(26), 8388–8395. <https://doi.org/10.1523/JNEUROSCI.0717-09.2009>
- Frith, C. D., & Frith, U. (2012). Mechanisms of Social Cognition. *Annual Review of Psychology*, 63(1), 287–313. <https://doi.org/10.1146/annurev-psych-120710-100449>
- Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., & Moser, M. B. (2004). Spatial representation in the entorhinal cortex. *Science*, 305(5688), 1258–1264. <https://doi.org/10.1126/science.1099901>
- Glasser, M. F., Smith, S. M., Marcus, D. S., Andersson, J. L. R., Auerbach, E. J., Behrens, T. E. J., ... Van Essen, D. C. (2016). The Human Connectome Project's neuroimaging approach. *Nature Neuroscience*, 19(9), 1175–1187. <https://doi.org/10.1038/nn.4361>
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B. B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052), 801–806. <https://doi.org/10.1038/nature03721>
- Hampton, A. N., Bossaerts, P., & O'Doherty, J. P. (2008). Neural correlates of mentalizing-related computations during strategic interactions in humans. *Proceedings of the National Academy of Sciences of the United States of America*, 105(18), 6741–6746. <https://doi.org/10.1073/pnas.0711099105>
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*.
- Hunt, L. T., Kolling, N., Soltani, A., Woolrich, M. W., Rushworth, M. F. S., & Behrens, T. E. J. (2012). Mechanisms underlying cortical activity during value-guided choice. *Nature Neuroscience*, 15(3), 470–476. <https://doi.org/10.1038/nn.3017>
- Jacobs, J., Weidemann, C. T., Miller, J. F., Solway, A., Burke, J. F., Wei, X.-X. X., ... Kahana, M. J. (2013). Direct recordings of grid-like neuronal activity in human spatial navigation. *Nature Neuroscience*, 16(9), 1188–1190. <https://doi.org/10.1038/nn.3466>
- Jenkinson, M., & Woolrich, M. (2004). Asymptotic T to Z and F to Z Statistic Transformations. Retrieved from <https://www.fmrib.ox.ac.uk/datasets/techrep/tr00mj1/tr00mj1/>
- Julian, J. B., Keinath, A. T., Frazzetta, G., & Epstein, R. A. (2018). Human entorhinal cortex represents visual space using a boundary-anchored grid. *Nature Neuroscience*, 21(2), 191–194. <https://doi.org/10.1038/s41593-017-0049-1>
- Kendall, M. G. (1938). A New Measure of Rank Correlation. *Biometrika*, 30(1/2), 81. <https://doi.org/10.2307/2332226>
- Knudsen, E. B., & Wallis, J. D. (2020). Closed-Loop Theta Stimulation in the Orbitofrontal Cortex Prevents Reward-Based Learning. *Neuron*, 106(3), 537–547.e4. <https://doi.org/10.1016/j.neuron.2020.02.003>
- Kolling, N., Behrens, T. E. J., Mars, R. B., & Rushworth, M. F. S. (2012). Neural Mechanisms of Foraging. *Science*, 336(6077), 95–98. <https://doi.org/10.1126/science.1216930>
- Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2(NOV), 4. <https://doi.org/10.3389/neuro.06.004.2008>
- Kriete, T., Noelle, D. C., Cohen, J. D., & O'Reilly, R. C. (2013). Indirection and symbol-like processing in the prefrontal cortex and basal ganglia. *Proceedings of the National Academy of Sciences*, 110(41), 16390–16395.

- <https://doi.org/10.1073/pnas.1303547110>
- Kumaran, D., Melo, H. L., & Duzel, E. (2012). The Emergence and Representation of Knowledge about Social and Nonsocial Hierarchies. *Neuron*, 76(3), 653–666. <https://doi.org/10.1016/j.neuron.2012.09.035>
- Lebreton, M., Abitbol, R., Daunizeau, J., & Pessiglione, M. (2015). Automatic integration of confidence in the brain valuation signal. *Nature Neuroscience*, 18(8), 1159–1167. <https://doi.org/10.1038/nn.4064>
- Lim, S.-L., O'Doherty, J. P., & Rangel, A. (2011). The Decision Value Computations in the vmPFC and Striatum Use a Relative Value Code That is Guided by Visual Attention. *Journal of Neuroscience*, 31(37), 13214–13223. <https://doi.org/10.1523/jneurosci.1246-11.2011>
- Long, X., & Zhang, S.-J. (2018). A novel somatosensory spatial navigation system outside the hippocampal formation. *BioRxiv*, 473090. <https://doi.org/10.1101/473090>
- Maidenbaum, S., Miller, J., Stein, J. M., & Jacobs, J. (2018). Grid-like hexadirectional modulation of human entorhinal theta oscillations. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.1805007115>
- Mars, R. B., Sallet, J., Schuffelgen, U., Jbabdi, S., Toni, I., & Rushworth, M. F. S. S. (2012). Connectivity-Based Subdivisions of the Human Right “Temporoparietal Junction Area”: Evidence for Different Areas Participating in Different Cortical Networks. *Cerebral Cortex*, 22(8), 1894–1903. <https://doi.org/10.1093/cercor/bhr268>
- Mikl, M., Mareček, R., Hlušík, P., Pavlicová, M., Drastich, A., Chlebus, P., ... Krupa, P. (2008). Effects of spatial smoothing on fMRI group inferences. *Magnetic Resonance Imaging*, 26(4), 490–503. <https://doi.org/10.1016/j.mri.2007.08.006>
- Nau, M., Navarro Schröder, T., Bellmund, J. L. S., & Doeller, C. F. (2018). Hexadirectional coding of visual space in human entorhinal cortex. *Nature Neuroscience*, 21(2), 188–190. <https://doi.org/10.1038/s41593-017-0050-8>
- Neubert, F.-X., Mars, R. B., Sallet, J., & Rushworth, M. F. S. (2015). Connectivity reveals relationship of brain areas for reward-guided learning and decision making in human and monkey frontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 112(20), 1–10. <https://doi.org/10.1073/pnas.1410767112>
- Nicolle, A., Klein-Flügge, M. C., Hunt, L. T., Vlaev, I., Dolan, R. J., & Behrens, T. E. J. (2012). An Agent Independent Axis for Executed and Modeled Choice in Medial Prefrontal Cortex. *Neuron*, 75(6), 1114–1121. <https://doi.org/10.1016/j.neuron.2012.07.023>
- Noonan, M. P., Mars, R. B., & Rushworth, M. F. S. (2011). Distinct roles of three frontal cortical areas in reward-guided behavior. *Journal of Neuroscience*, 31(40), 14399–14412. <https://doi.org/10.1523/JNEUROSCI.6456-10.2011>
- O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Clarendon Press.
- Park, S. A., Sestito, M., Boorman, E. D., & Dreher, J. C. (2019). Neural computations underlying strategic social decision-making in groups. *Nature Communications*, 10(1), 5287. <https://doi.org/10.1038/s41467-019-12937-5>
- Piva, M., Velnoskey, K., Jia, R., Nair, A., Levy, I., & Chang, S. W. C. (2019). The dorsomedial prefrontal cortex computes task-invariant relative subjective value for self and other. *ELife*, 8. <https://doi.org/10.7554/eLife.44939.001>
- Platt, M. L., Seyfarth, R. M., & Cheney, D. L. (2016). Adaptations for social cognition in the primate brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1687). <https://doi.org/10.1098/rstb.2015.0096>
- Preston, A. R., & Eichenbaum, H. (2013, September 9). Interplay of hippocampus and prefrontal cortex in memory. *Current Biology*. Cell Press. <https://doi.org/10.1016/j.cub.2013.05.041>
- Schuck, N. W., Cai, M. B., Wilson, R. C., & Niv, Y. (2016). Human Orbitofrontal Cortex Represents a Cognitive Map of State Space. *Neuron*, 91(6), 1402–1412. <https://doi.org/10.1016/J.NEURON.2016.08.019>
- Smith, S. M., & Nichols, T. E. (2009). Threshold-free cluster enhancement: Addressing problems of smoothing, threshold dependence and localisation in cluster inference.

- NeuroImage*, 44(1), 83–98. <https://doi.org/10.1016/j.neuroimage.2008.03.061>
- Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, 20(11), 1643–1653. <https://doi.org/10.1038/nn.4650>
- Staudigl, T., Leszczynski, M., Jacobs, J., Sheth, S. A., Schroeder, C. E., Jensen, O., & Doeller, C. F. (2018). Hexadirectional Modulation of High-Frequency Electrophysiological Activity in the Human Anterior Medial Temporal Lobe Maps Visual Space. *Current Biology*. <https://doi.org/10.1016/j.cub.2018.09.035>
- Strohming, N., Gray, K., Chituc, V., Heffner, J., Schein, C., & Heagins, T. B. (2016). The MR2: A multi-racial, mega-resolution database of facial stimuli. *Behavior Research Methods*, 48(3), 1197–1204. <https://doi.org/10.3758/s13428-015-0641-9>
- Sutton, R. S., & Barto, A. G. (1998). Reinforcement Learning: An Introduction, 322. <https://doi.org/10.1109/TNN.1998.712192>
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55(4), 189–208. <https://doi.org/10.1037/h0061626>
- Walther, A., Nili, H., Ejaz, N., Alink, A., Kriegeskorte, N., & Diedrichsen, J. (2016). Reliability of dissimilarity measures for multi-voxel pattern analysis. *NeuroImage*, 137, 188–200. <https://doi.org/10.1016/j.neuroimage.2015.12.012>
- Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., ... Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, 21(6), 860–868. <https://doi.org/10.1038/s41593-018-0147-8>
- Weiskopf, N., Hutton, C., Josephs, O., & Deichmann, R. (2006). Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: A whole-brain analysis at 3 T and 1.5 T. *NeuroImage*, 33(2), 493–504. <https://doi.org/10.1016/j.neuroimage.2006.07.029>
- Whittington, J. C., Muller, T. H., Mark, S., Chen, G., Barry, C., Burgess, N., & Behrens, T. E. (2019). The Tolman-Eichenbaum Machine: Unifying space and relational memory through generalisation in the hippocampal formation. *BioRxiv*, 770495. <https://doi.org/10.1101/770495>
- Witter, M. P., Wouterlood, F. G., Naber, P. A., & Van Haeften, T. (2000). Anatomical organization of the parahippocampal-hippocampal network. *Annals of the New York Academy of Sciences*, 911(1), 1–24. <https://doi.org/10.1111/j.1749-6632.2000.tb06716.x>
- Wittmann, M. K., Lockwood, P. L., & Rushworth, M. F. S. (2018). Neural Mechanisms of Social Cognition in Primates. *Annual Review of Neuroscience*. <https://doi.org/10.1146/annurev-neuro-080317-061450>
- Yamada, H., Louie, K., Tymula, A., & Glimcher, P. W. (2018). Free choice shapes normalized value signals in medial orbitofrontal cortex. *Nature Communications*. <https://doi.org/10.1038/s41467-017-02614-w>
- Yushkevich, P. A., Amaral, R. S. C., Augustinack, J. C., Bender, A. R., Bernstein, J. D., Boccardi, M., ... Zeineh, M. M. (2015). Quantitative comparison of 21 protocols for labeling hippocampal subfields and parahippocampal subregions in in vivo MRI: Towards a harmonized segmentation protocol. *NeuroImage*, 111, 526–541. <https://doi.org/10.1016/j.neuroimage.2015.01.004>
- Zhou, J., Montesinos-Cartagena, M., Wikenheiser, A. M., Gardner, M. P. H., Niv, Y., & Schoenbaum, G. (2019). Complementary Task Structure Representations in Hippocampus and Orbitofrontal Cortex during an Odor Sequence Task. *Current Biology*, 29(20), 3402–3409.e3. <https://doi.org/10.1016/j.cub.2019.08.040>
- Zilles, K., & Amunts, K. (2010). Centenary of Brodmann's map conception and fate. *Nature Reviews Neuroscience*, 11(2), 139–145. <https://doi.org/10.1038/nrn2776>

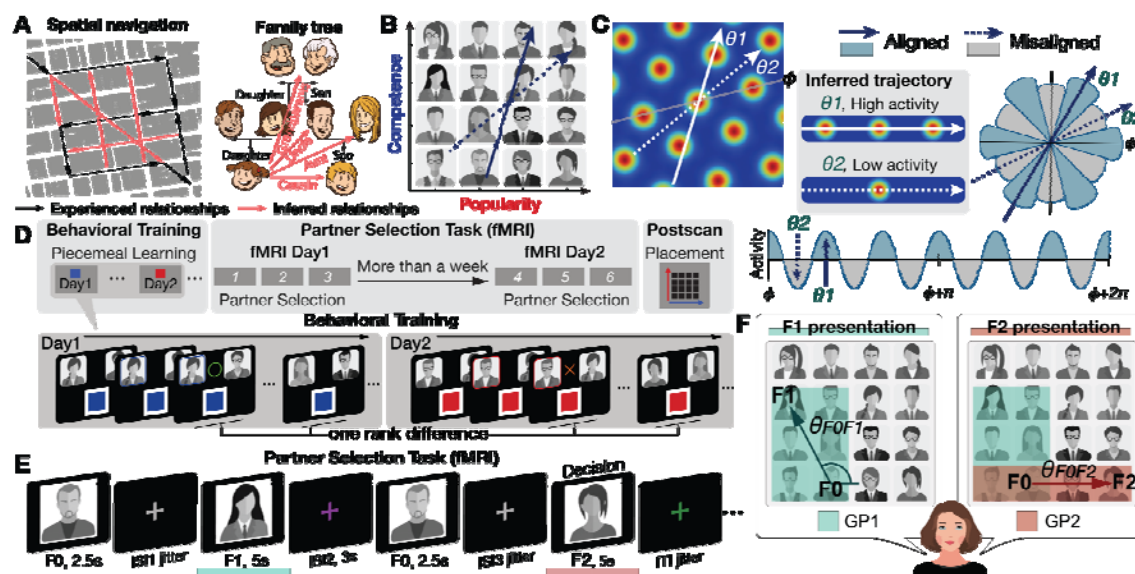


Figure 1. Behavioral training and experimental task. **A.** A cognitive map of spatial (left) and non-spatial (right) relational structures allows for new direct routes or relationships to be inferred from directly experienced relationships, dramatically accelerating learning and decision capabilities through generalization. **B.** Participants learned the rank of 16 individuals organized into a 2-D social hierarchy defined by competence and popularity. Participants never saw neither 1-D or 2-D social hierarchies, but they could learn it piecemeal from dyadic comparisons in one dimension at a time during behavioral training. We hypothesized neural activity would be modulated hexadirectionally by the inferred trajectories over the 2-D social space, as predicted by a hexagonal grid organization. **C.** The inferred trajectories can be categorized as aligned and misaligned with the mean grid orientation, θ , which is different for each participant. θ_1 and θ_2 show examples of aligned and misaligned trajectories, respectively. Greater activity is predicted when the inferred trajectory is aligned compared to misaligned because it passes over more grid fields, which generates hexadirectional grid-like modulation. **D.** Behavioral training procedure. On day 1 and day 2 of behavioral training, participants learned ranks of 16 individuals (face stimuli) in each of two dimensions (competence or popularity) through binary decisions about the higher rank individual in a pair who differed by only one rank level in a given dimension, with each dimension learned on a different day. Within a day, the order of pairs compared was further randomized. After behavioral training, participants performed 3 blocks of the "partner selection task" twice during fMRI scanning, with a gap between sessions of at least one week. After the second session, participants were asked to place individuals according to their believed combined rank in a 2-D space for the first time (placement task; **Fig. S1F**). **E.** Illustration of a trial of the partner selection task during fMRI. Participants were asked to make a binary decision by choosing a better business partner for a given individual (F0) between two (F1 and F2). The better partner is determined by the "growth potential (GP)" that each pair could expect from their collaboration. Participants could compute the GP of the F0 and F1 pair when F1 is shown and GP of the F0 and F2 pair when F2 is shown. Participants were subsequently asked to make a decision during F2 presentation. No feedback was given. **F.** To compute the GP of a pair, participants were instructed the GP corresponds to the higher "rank of the pair" in each dimension. Participants were further asked to weigh the 'rank of the pair' in both dimensions equally. Therefore, the GP corresponds to the area drawn by the higher rank of the two people in each dimension in the 2-D hierarchy (GP_{F0F1} (green rectangle) > GP_{F0F2} (red rectangle); F1 is the better partner for F0 in this example). We hypothesized people would infer direct trajectories over the mentally reconstructed 2-D space between the positions of F0 and each potential partner, F1 and F2, to compute the GP for each collaboration. We searched for neural evidence for

hexadirectional modulation of inferred trajectories though the reconstructed cognitive map (F0F1 at the time of F1 presentation and F0F2 at the time of F2 presentation).

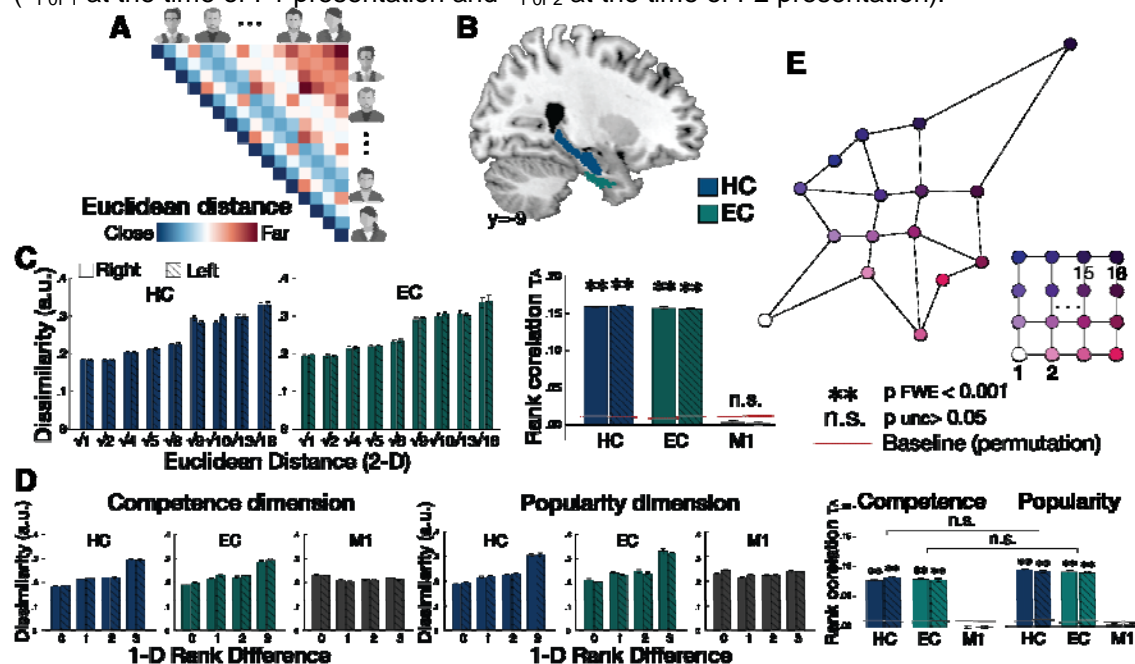


Figure 2. Representational similarity analysis (RSA). **A.** Left: Model representational dissimilarity matrix (RDM) that models distances computed from the pairwise Euclidean distances between individuals on the true 2-D social hierarchy. **B.** We examine to what extent the model RDM explained the pattern dissimilarity between voxels in anatomically defined *a priori* regions of interest (ROIs) in bilateral hippocampus (HC) and entorhinal cortex (EC). The neural RDM was computed in the ROIs based on the pairwise Mahalanobis distance in the multi-voxel activity patterns evoked when face stimuli were presented. **C.** The dissimilarity between activity patterns estimated in bilateral HC and EC increases in proportion to the true pairwise Euclidean distance between individuals in the 2-D abstract social space. The rank correlation (Kendall's τ_A) shows robust effects of Euclidean distance on the pattern dissimilarity estimated in the HC and EC compared to the permuted baseline (1000 iterations), but not in a control region, the primary motor cortex (M1) (**, $p_{FWE} < 0.001$ after correction for the number of bilateral ROIs ($n=4$) with the Bonferroni-Holm method; n.s., $p > 0.05$, uncorrected). **D.** The dissimilarity between activity patterns estimated in bilateral HC and EC increases not only with the 1-D rank distance in competence dimension (left) but also the 1-D rank distance in popularity dimension (middle). Right: The rank correlation of one dimension was not different from that of the other dimension ($p > 0.05$), suggesting that the HC-EC system utilizes a 2-D relational cognitive map to represent the social hierarchy. **E.** Visualization of the group representation of the social hierarchy in a 2-D space using multidimensional scaling (MDS, left) on the neural activity extracted from the HC ROIs. The distances between the estimated individual locations in a 2-D space are significantly correlated with the pairwise Euclidean distances in the true 4x4 social hierarchy structure (right) ($r = 0.80$, $p < 0.001$). Note that there were considerably fewer trials for rank positions 1 and 16, making their estimates less reliable.

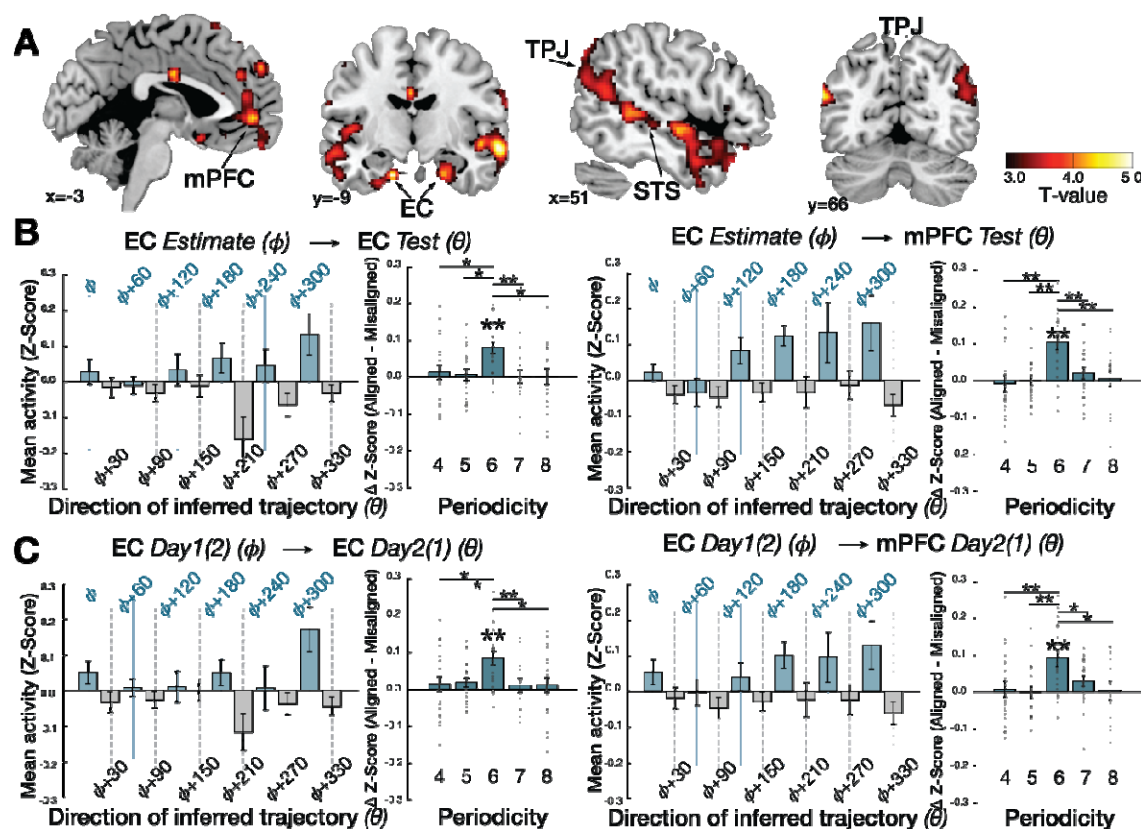


Figure 3. Hexagonal modulation for inferred trajectories. **A.** Whole-brain parametric analysis showing hexagonal grid-like representation of inferred trajectories in alignment with the mean EC grid orientation at the time of F1 and F2 presentations. Significant effects are shown in EC (peak MNI coordinates, [x,y,z]=[22,-10,-28], t=4.11), mPFC ([-6,48,-4]; t=4.72), STS ([50,-40,4], t=4.05 for right; [-60,-24,-6], t=4.29 for left), TPJ ([46,-58,20], t=3.67 for right; [-56,-68,24], t=5.71 for left) (all $p_{TFCE} < 0.05$, whole-brain cluster corrected using threshold-free cluster enhancement (TFCE) (Smith & Nichols, 2009)). For visualization purposes, the whole-brain maps are thresholded at $p < 0.005$ uncorrected. **B.** Six-fold modulation signals in the EC (left panel) and mPFC (right panel) ROIs aligned to the grid orientation in EC. The grid orientation was estimated from separate fMRI sessions acquired from the same day. The mean (\pm SE) z-scored activity is plotted separately for aligned (teal) and misaligned (gray) trajectories categorized into 12 equal bins of 30° according to the direction of inferred trajectories. The mean activity difference between aligned and misaligned trajectories was larger than zero for six-fold ($p < 0.01$) but not for the other control periodicities (four-, five-, seven- and eight-fold; all $p > 0.05$), and the activity difference is greater for six-fold compared to the other periodicities ($p < 0.05$). **C.** Cross-day consistency of the grid orientation in EC (See Fig. S4B for whole brain analysis). The activity in EC and mPFC shows hexadirectional modulations for the inferred trajectories in alignment with the grid orientation in EC estimated from separate sessions acquired from a different day more than a week apart. This effect is also specific to the six-fold ($p < 0.01$) periodicity (all $p > 0.05$), suggesting that the mean grid angle in EC is consistent between sessions more than a week apart. **, $p < 0.01$, *, $p < 0.05$.

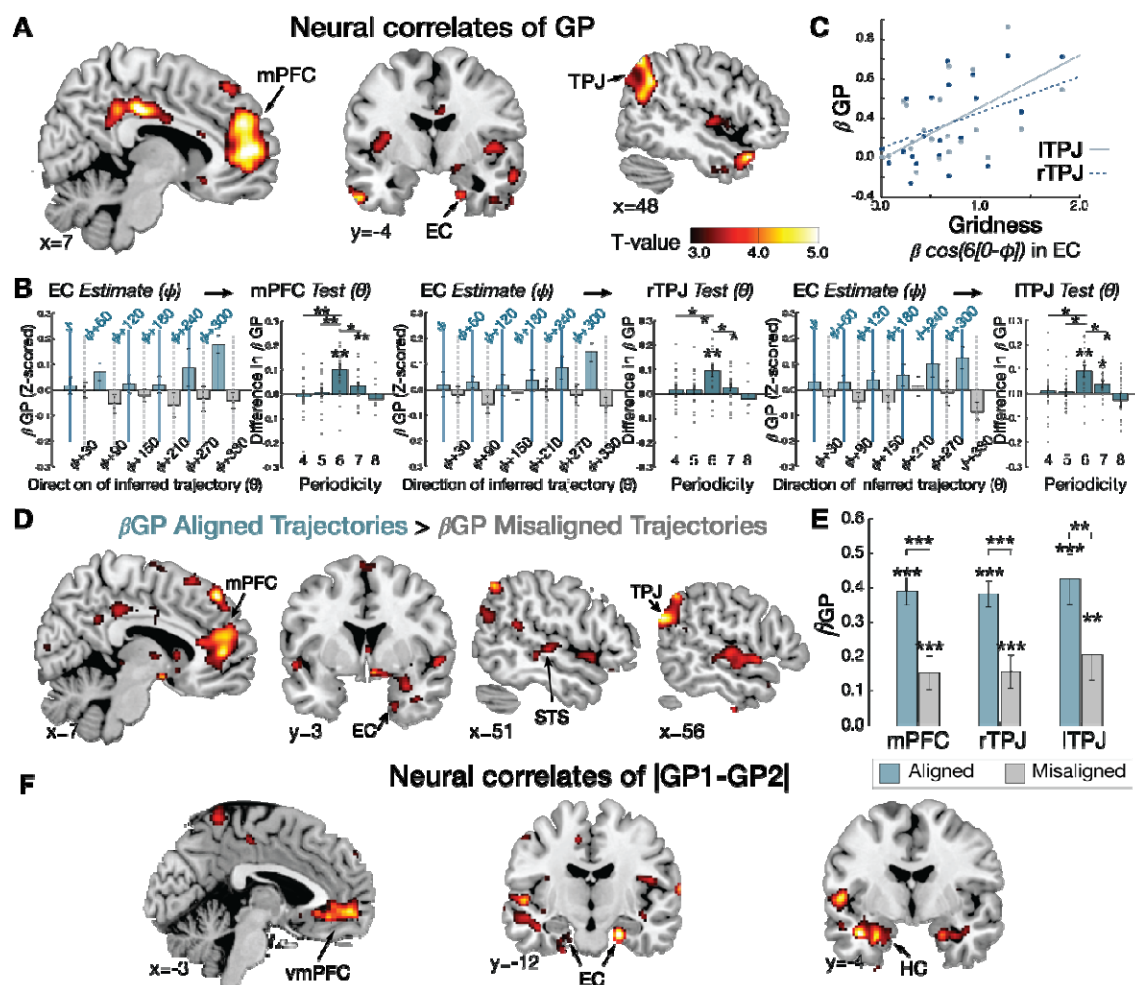


Figure 4. Growth potential and value comparison. **A.** Whole-brain map showing effects of the “growth potential” (GP) at the time of F1 and F2 presentation in a network of brain regions including the mPFC (peak MNI coordinates, [x,y,z]=[10,52,6], $t=5.55$) and bilateral TPJ ([54,-56,34], $t=5.45$) for right; [-54,-60,28], $t=5.18$) for left ($p_{TFCE} < 0.05$ whole brain TFCE correction). The EC ([x,y,z]=[20,-4,-32], $t=3.88$) also showed GP effects ($p_{TFCE} < 0.05$ TFCE-corrected in an anatomically defined ROI). **B.** Mean (\pm SE) effects of GP (z-scored β) in mPFC and bilateral TPJ are modulated by the grid alignment of the inferred trajectories, , with six-fold periodicity aligned with the EC grid orientation, . Control analyses confirmed specificity for this six-fold periodicity ($p < 0.01$) over control periodicities (effect of six-fold compared to other control periodicities, all $p < 0.05$). **C.** Those participants with greater hexagonal modulation in EC show greater encoding of GP in independently defined bilateral TPJ ROIs (all $p < 0.05$). **D.** Whole-brain map contrasting the effects of GP (β GP) for trajectories aligned with the EC grid orientation, , compared to the misaligned trajectories (β GP aligned > β GP misaligned: $p_{TFCE} < 0.05$ whole brain TFCE correction). Contrasts are shown in mPFC, EC, STS, and TPJ. **E.** Mean (\pm SE) GP effects in anatomically mPFC, and bilateral TPJ (β GP). In all three areas, effects were greater for aligned compared to misaligned trajectories (all $p < 0.01$), although the GP effects of both trajectories were significant (all $p < 0.01$). **F.** Neural correlates of the relative decision value, |GP1-GP2|, during decision making, in vmPFC ([-8,54,-6]; $t=5.07$), HC ([-28,-4,-26]; $t=4.53$) ($p_{TFCE} < 0.05$), and bilateral EC/HC ([22,-12,-26], $t=5.10$ for right; [-26,-16,-32], $t=4.00$ for left; $p_{TFCE} < 0.05$ TFCE corrected in *a priori* ROIs). For visualization purposes, the whole-brain maps are thresholded at $p < 0.005$ uncorrected. ***, $p < 0.005$; **, $p < 0.01$; *, $p < 0.05$.

METHODS

Participants

A total of 25 participants (13 females, age range: 18–25, normal or corrected to normal vision) were recruited for this study via the University of California, Davis online recruitment system. Using six motion parameters to assess motion in each participant, we excluded the data of four participants who had movement in the x, y, and/or z directions greater than the voxel size (3 mm) relative to the middle time point across each scanning block and who had total displacement of greater than 3 mm from the beginning to the end of the scanning session. In total, 21 participants' data were included in the analysis (10 female, mean age: 21.14 ± 0.52 , standard error mean (\pm SE)). The study was approved by the local ethics committee, all relevant ethical regulations were followed, and participants gave written consent before the experiment.

Participants received course credit as compensation for participating in day 1 and day 2 training. Participants received monetary compensation for their participation from day 3 training. In addition to the compensation, participants have been told that extra monetary rewards will be earned if they pass the training (see the criteria below in behavioral training) and these extra rewards would be paid if they completed the fMRI task.

Social hierarchies

During behavioral training, participants were asked to learn social hierarchies between 16 entrepreneurs presented as face stimuli which were constructed in two independent ability dimensions – popularity and competence. Each dimension had four levels of rank. Four individuals were allocated at the same rank at each level. The rank of an individual in one dimension was not related to his/her rank in the other dimension and each had a unique combination of popularity and competence ranks. This allowed us to create a 4x4 social hierarchy structure comprising two social hierarchy dimensions (**Fig. 1B**).

It is important to note that the true structure of social hierarchy (not only the two-dimensional (2-D) social hierarchy but also each one-dimensional (1-D) social hierarchy) was never shown to participants. Moreover, participants were never given any information implying the structure of the social hierarchies, including the total number of ranks in each dimension and the number of individuals allocated to the same rank.

Stimuli

The stimuli consisted of 16 grayscale photographic images of faces (Strohming et al., 2016). All of the images were adjusted to have the same mean grayscale value. To ensure that visual features of a face (including gender, race, and age) were not associated with the rank of the entrepreneur, we prepared eight sets of 16 face stimuli selected from a pool of 42 faces. Each of the stimuli sets comprised 16 faces. A set of stimuli among eight was randomly assigned across participants. Stimuli were presented to participants through a mirror mounted on the head coil. The face stimuli presented in this paper are license-free images for display purposes.

Behavioral training

The 16 face stimuli were introduced as entrepreneurs. Each day of the behavioral training comprised 'Learn' phases and 'Test' phases. During a Learn phase of behavioral training, participants learned the relative rank between entrepreneurs (face stimuli) in one of the two independent dimensions through feedback-based binary comparisons. Specifically,

participants were asked to learn which individuals were more capable of attracting crowdfunds (labelled popularity) and which individuals had higher technical proficiency (labelled competence). Participants only learned about individuals who differed by one rank level and on one dimension at a time. Participants learned the relative status of all possible one rank difference pairs with feedback in random order. It is Important to note that the relative ranks in each dimension were learned on a separate day (e.g. popularity on day 1 and competence on day 2 or vice versa). The day 2 training occurred 2 days after day 1 training. During day 2 training, participants learned the hierarchies on the other, different dimension (**Fig. 1D and Fig. S1**).

Participants were never shown the 1-D or 2-D structure of the social hierarchy, never asked to combine the two social hierarchy dimensions during training, and never asked to solve the task spatially. However, participants could theoretically reconstruct the social hierarchy through transitive inferences not only on two one-dimensional structures but also on a two-dimensional structure by integrating the two dimensions learned on different days (as shown in the right panels in **Fig. S1**). During training, participants were not been informed about the fMRI task (Partner selection task). Therefore, participants were not able to anticipate having to build a representation of the combined social hierarchy for future use for novel inferences. During behavioral training, each of the 16 face stimuli had been exposed to participants the same number of times.

Training day 1

Participants learned the hierarchy of 16 individuals in either the competence or popularity dimension. On each trial during the Learn phase, two face stimuli of entrepreneurs were presented with a color cue (red or blue square) which indicated the task-relevant dimension. Participants were asked to choose the one who is superior to the other (the higher rank individual) in the given dimension. They received feedback at the end of every trial. Before the training, participants were informed that two entrepreneurs presented in a given trial during the Learn phase differed by only one rank on the given dimension (**Fig. S1A**). The first learned hierarchy and the color associated with each dimension were counterbalanced across participants.

During the Test phase, which was followed by the Learn phase, participants' knowledge about social hierarchy in the learned dimension was tested. On each trial during the Test phase, similar to the Learn phase, participants chose one of two face stimuli who is superior to the other in the given social hierarchy. Participants did not receive any feedback for their choices during the Test phase. Before the training, participants were informed that two entrepreneurs presented in the Test phase would have one or more rank differences. To make a correct decision, participants needed to use transitive inferences to infer the relative status of two individuals who have never been compared during learning (**Fig. S1A**).

Training day 2

Using the same method through which the participants learned the social hierarchy in one dimension during day 1 training, participants learned the hierarchy of the same 16 entrepreneurs in the other unlearned dimension. For example, participants who learned about the hierarchy in competence dimension on day 1 learned about the hierarchy in popularity dimension on day 2 (see **Fig. S1B**).

At the end of day 2 training, participants were asked to perform an additional test which we called the "Test 2" phase. During the Test 2 phase, we tested participants' ability to make flexible inferences while the task-relevant dimension was interleaved across trials.

As with the previous Test phase, on each trial of the Test 2 phase, participants were asked to choose the superior face between two entrepreneurs who have one or more rank differences in the given dimension, and no feedback was given for their decisions. However, in the Test 2 phase, the relevant dimension for decision making was randomized across trials. Only the participants who were able to reach more than 80% accuracy in the Test 2 phase were invited to continue to the day 3 training (**Fig. S1B**).

Day 3 and after training

A. Flexible inferences in intermixed behavioral contexts (Test 2)

The day 3 training started with another Test 2 phase in which participants' knowledge about social hierarchy was tested with the task-relevant dimension randomized across trials. As before, participants were asked to infer the relative rank between two individuals with one or more rank differences without feedback. From the day 3 training, on each trial of Test 2 phase, stimuli were presented sequentially in the order of the conditional cue (red or blue square) indicating the task-relevant dimension, a face stimulus (F1), and another face stimulus (F2) (**Fig. S1C**). A fixation cross was shown for the inter-stimulus-intervals and inter-trial intervals. During F2 presentation (2s), participants were asked to indicate the higher rank individual between F1 and F2 in the given dimension. During the Test 2 phase, participants were presented all possible pairs of entrepreneurs twice by altering the position of F1 and F2 on interleaved dimensions, and they did not get feedback to their decisions. Only the participants who could make accurate inferences at 95% or more in both dimensions finished the training and were given the introduction about the partner selection task.

B. Conditional Learn phase

There were additional Learn phases for the participants whose accuracy was below 95% in one of two dimensions for the Test 2 phase on day 3 training. This additional Learn phase comprised the two Learn phases given on the day 1 and day 2 training. That is, participants were asked to choose the superior face between two entrepreneurs who differed by one rank on the given dimension while getting feedback to their decisions. Participants learned the hierarchy in one dimension at a time during each of two blocks in the Learn phase (**Fig. S1C**).

After the Learn phase, participants were asked to repeat the Test 2 phase of day 3 training. It could potentially take multiple days of training until the threshold was reached (>95% in both dimensions) (**Fig. S1C**). This high performance threshold was selected because we reasoned it would be necessary to accurately measure hypothesized grid angles for planned fMRI analyses.

Behavioral training of the “partner selection” task

Participants who successfully learned the social hierarchy in both dimensions independently subsequently learned the partner selection task for the first time. In each trial of partner selection task, participants were asked to choose the better business partner for a given entrepreneur between two candidate partners while considering their relative hierarchy in both social hierarchy dimensions. Therefore, each decision during the fMRI task required novel inferences to be made by combining the two social hierarchy dimensions learned separately. See the following section of “Partner selection task” for details.

The behavioral training of the partner selection task had two purposes. First, it was to ensure participants understood the novel decision-making task in which they had never been

exposed. Second, it was to ensure participants did not use alternative heuristics for decision-making by providing test trials where the potential heuristics failed (See the Partner selection task for details). To do this, after the instructions were given, participants were asked to perform 24 trials of the partner selection task without time constraints. To prevent learning, we did not give any feedback to their decisions during the training of the partner selection task. They were only informed whether they made all correct selections or not at the end of a block which comprised of 24 trials. We generated three sets of 24 trials for the purpose of behavioral training. Those trials were not included in the fMRI experiments to make sure participants made novel decisions for the fMRI experiment. Participants who failed to make all correct inferences in 24 trials were further asked to read again the instructions and to explain the instruction to the experimenter to ensure that they understood the task. Participants were asked to complete the same set of 24 trials in a randomized order until they made all correct decisions. After reaching 100% accuracy, participants were asked to complete the same set of 24 trials in a randomized order with time constraints (5 s), like in the fMRI experiment. To ensure that participants were not simply relying on remembered correct answers of the selected set were able to generalize their knowledge to make novel inferences, we confirmed that all participants could also make correct decisions for a novel set of 24 trials of the partner selection task. Finally, using this procedure, we recruited participants who were reach an accuracy of 95% or more in both dimensions for more than two trial sets (48 trials total) to participate in the fMRI experiment.

Partner selection task (fMRI experiments)

Participants were given the following instructions. The goal of the partner selection task was to find a better business partner for a given individual (F0) between two candidate partners (F1 and F2). F0, F1, and F2 were indicated as face stimuli which were selected from the 16 entrepreneurs that had been presented to participants during behavioral training. To find a better business partner for F0, participants need to compute the 'growth potential' (GP) of each potential collaboration. The GP indicates the level of benefits that F0 could expect from the potential collaboration with each of two partners, F1 and F2. Participants were subsequently asked to choose between F1 and F2 according to the GP of each of two pairs – GP1: GP for F0-F1 pair; GP2: GP for F0-F2 pair.

To estimate GP accurately, first, participants were asked to compute the 'rank of the pair' in each dimension which was determined by the rank of the entrepreneur who had a higher rank (where the lowest rank was one). For example, if F0 was at a higher rank than F1 in the competence dimension while F1 was at a higher rank than F0 in the popularity dimension, the rank of the F0-F1 pair would be the rank of F0 in competence dimension and rank of F1 in popularity dimension. Second, participants were asked to weight the 'rank of the pair' in both dimensions equally to compute the GP. For example, when the ranks of F0 and F1 in the competence and popularity dimensions are $[C_{F0}, P_{F0}]$ and $[C_{F1}, P_{F1}]$, respectively, the 'rank of the pair' would be $\max(C_{F0}, C_{F1})$ in the competent dimension and $\max(P_{F0}, P_{F1})$ in the popularity dimension. And the GP1 would be determined by multiplication of two 'ranks of the pair'. That is, $\max(C_{F0}, C_{F1}) \times \max(P_{F0}, P_{F1})$. Last, participants were asked to make a binary choice between F1 and F2 to find a better partner for F0 by selecting the pair whose GP was greater than the other pair: choose F1 if $GP1 > GP2$; choose F2 if $GP2 > GP1$. During the partner selection task, participants did not receive any feedback for their decisions. Participants were notified that they would receive extra rewards in proportion to their overall performance, if their overall accuracy was greater than 90% at the end of the experiments (after the second-day scan).

As **Fig. 1E** illustrates, in each trial of the partner selection task, participants were presented with two pairs of face stimuli sequentially in the order of F0, F1, F0, and F2. The stimuli, F0 were presented for 2.5 s and the stimuli, F1 and F2 were presented for 5s. In each trial, the inter-stimulus fixation was presented between F0 and F1 and between F0 and F2 for the same duration (a white cross, 2~5 s jittered). F0 was presented before each partner to anchor hypothesized direct vectors to the partner. The F1 stimuli were followed by a 3 s an inter-stimuli fixation (a purple cross). Participants were asked to make a binary decision to indicate a better business partner for F0 between F1 and F2 during the F2 presentation by pressing a button. The inter-trial fixation was presented at the end of each trial (green cross, 1~4 s jittered). During the fMRI experiment, the color of the fixation cross on the middle of the screen informed participants the progress of the trial. No feedback was given to the participants. Trials that participants did not respond to within the allotted time were shown later after a random number of trials to make sure that all subjects made responses to every trial.

fMRI experimental procedure

Subjects participated in two days of scanning sessions more than a week apart (**Fig. 1C**). The average time gap between the two sessions was 8.95 ± 0.81 days (\pm SE) (in the range of 7 to 19 days). On each day of the fMRI experiment session, participants performed three blocks of the partner selection task. A block consisted of 48 trials. Participants, therefore, needed to compute the GP of 576 pairs in total (48 trials \times 2 pairs per trial \times 3 blocks \times 2 days) to make accurate decisions for the partner selection task. The pairs presented during the first-day scan were re-presented for the second-day scan in the reverse order. That is, the F0 and F1 pairs in the first-day scan were presented as the F0 and F2 pairs in the second-day scan and vice versa. This allowed us to test whether participants made correct decisions consistently as well as whether their brain consistently represented the direction of inferred trajectories (θ) over the abstract 2-D social hierarchy space across sessions acquired more than a week apart. **Fig. S1D** shows the frequency of the inferred trajectories for the fMRI experiment categorized into 12 equal bins of 30° according to their direction (θ), which indicates that the 2-D social hierarchy space was well sampled. While we randomized the trial order in each block, we also changed the order of blocks between the first day scan and the second day scans. That is, the trials tested in the blocks 1, 2, and 3 on the first day scan were tested in the blocks 5, 6, and 4, respectively on the second-day scan, which allowed us to generate unique comparisons for cross-validation across blocks (see the fMRI analysis).

Post-scanning debriefing

After all the scans finished (on the second-day scan), we asked participants to describe the mental images that they had used to represent the interrelated 16 face stimuli and their decision strategy used to perform the partner selection task. All participants mentioned that when they see a face stimulus, they did not recall the rank in one dimension at a time, but they were able to recall the ranks in both dimensions simultaneously. Among 21 participants whose data were analyzed, only two participants reported the 4×4 social hierarchy structure in a 2-D space they used to represent the relationship. Other participants did not report being aware that this task can be solved spatially using an explicit 2-D representation of the social hierarchy.

Post-scanning placement task

After participants reported their strategy, we asked them to perform a placement task. We asked participants to drag-and-drop each of the 16 faces of entrepreneurs to place on a 2-D space defined by the two social hierarchy dimensions according to their ranks. Though most of the participants reported that they did not use a spatial representation to perform the task during post-scanning debriefing, all participants were successfully able to place the faces according to their social hierarchy ranks and reconstruct the social hierarchies in a 4x4 grid.

For visualization purposes, we rescaled the map of the faces placed by each participant by normalizing the longest pairwise distance between faces on the X coordinate and that on the Y coordinate to one (**Fig. S1F**). Note that after rescaling, the location of faces can appear more precise at the boundary compared to the faces in the middle of the map. However, we confirmed that all participants reported an equidistant interval between any pair of neighboring entrepreneurs who differed by one rank in either the competence or popularity dimension.

Behavioral analyses

We examined whether participants made a model-based decision using the learned structure of the social hierarchy instead of a model-free decision using heuristics as an alternative decision strategy. Specifically, we examined whether the decision of participants was better explained by the difference between GPs of two pairs (GP1-GP2), which requires participants to assemble the ranks of individuals in both dimensions compared with the difference in ranks or values that participants assigned to each individual, which does not require the internal representation of the social hierarchy to guide a decision.

As the first alternative way to make decision, we tested whether participants used an overall rank assigned to each individual in a combined dimension rather than ranks of each of two dimensions and chose the one who had greater overall rank between F1 and F2. That is, this model predicted that participants choose F1 if $[P_{F1} \times C_{F1}] > [P_{F2} \times C_{F2}]$, and otherwise F2. This model differs from the GP-based decision in that the decision does not consider the rank of F0. As the second alternative model, we tested whether participants only use the ranks of F1 and F2 in one of two dimensions in which the rank of F0 was relatively deficient. For example, the decision of participants might only depend on the ranks of F1 and F2 in the popularity dimension when F0 was at a lower rank in the popularity dimension than their rank in the competence dimension. Formally, this model predicts that participants choose F1 if $C_{F1} > C_{F2}$ given that $P_{F0} > C_{F0}$ or if $P_{F1} > P_{F2}$ given that $C_{F0} > P_{F0}$, otherwise they choose F2. If participants use this heuristic, therefore, participants could have made sequential value comparisons, rather than using the inferred trajectories.

To test these alternative hypotheses, we included trials in which participants would fail to choose the correct partner if they adopted one of the alternatives. If a participant used the first heuristic, they would fail to be correct in 35 trials (24.3% of total trials). If a participant used the second heuristic, they would fail to be correct in 109 trials (75.7% of total trials). This allowed us to assess whether participants used the ranks of individuals in both dimensions to compute the expected GP, rather than using their ranks in only one dimension.

MRI data acquisition

We acquired T2-weighted functional images on a Siemens Skyra 3 Tesla scanner. We used gradient-echo-planar imaging (EPI) pulse sequence with a slice angle of 30° relative to the anterior-posterior commissure line, minimizing the signal loss in the orbitofrontal cortex region (Weiskopf, Hutton, Josephs, & Deichmann, 2006). We acquired 46 slices, 3mm thick

with the following parameters: repetition time (TR) = 1300 ms, echo time (TE) = 24 ms, flip angle = 67°, field of view (FoV) = 192mm, voxel size = 3 x 3 x 3 mm³. Contiguous slices were acquired in interleaved order. To correct for deformations, we also acquired a field map with dual echo-time images covering the whole brain, with the following parameters: TR = 630 ms, TE1 = 10 ms, TE2 = 12.46 ms, flip angle = 40°, FoV = 192mm, voxel size = 3 x 3 x 3 mm³. We acquired a T1-weighted structural image using a magnetization-prepared rapid gradient echo sequence (MPRAGE) with the following parameters: TR = 1810 ms, TE = 2.98 ms, flip angle = 7°, FoV = 256mm, voxel size = 1 x 1 x 1 mm³.

Pre-processing

The preprocessing of functional imaging data was performed using SPM12 (Wellcome Trust Centre for Neuroimaging). Images were corrected for slice timing, realigned to the first volume, and realigned to correct for motion using a six-parameter rigid body transformation. Inhomogeneities created using the phase of nonEPI gradient echo images at 2 echo times were coregistered with structural maps. Images were then spatially normalized by warping subject-specific images to the reference brain in MNI (Montreal Neurological Institute) coordinate space (2mm isotropic voxels). For the univariate analysis, images were smoothed using an 8-mm full-width at half maximum Gaussian kernel (Mikl et al., 2008).

fMRI representational similarity analysis (RSA)

We performed a multivariate analysis using RSA (Kriegeskorte et al., 2008) to test whether the *a priori* regions of interest (ROIs) in the brain reflected the structure of the social hierarchy (**Fig. 2B**). The ROIs were defined anatomically in the bilateral hippocampus (HC) (Yushkevich et al., 2015) and bilateral entorhinal cortex (EC) (Amunts et al., 2005; Zilles & Amunts, 2010). We also included additional ROIs in the bilateral primary motor cortex (M1) (Glasser et al., 2016) as control regions. Note that ROIs were defined independently from the current task.

By using RSA, we tested whether the brain integrates piecemeal learned relative status between entrepreneurs into a social hierarchy, and more importantly, whether the brain combines two social hierarchy dimensions, which were learned on different days, into a single 2-D representation. To address this question, we examined whether the level of pattern dissimilarity between the neural activity evoked by each of 16 entrepreneurs was explained by a function of pairwise Euclidean distances between entrepreneurs on the 2-D social space. This analysis allowed us to test whether the brain had increasingly similar patterns of neural activity for entrepreneurs who were increasingly close to each other in the true 2-D social space.

To test our hypotheses, we first estimated β coefficients when each of the individual faces was shown at the time of F0, F1 or F2 using a GLM with a 2s box-car function in each of our anatomically defined *a priori* ROIs. We then averaged the unsmoothed β maps when the same face among 16 faces was presented across F0, F1 and F2 presentations, allowing us to estimate the patterns of neural activity in each ROI. The reliability of the data was improved by applying multivariate noise normalization (Walther et al., 2016). Second, we quantified the representational similarity between the neural activity patterns elicited by face stimuli across three independent fMRI blocks acquired in the same day using the Mahalanobis distance. This pairwise dissimilarity further constructed a 16x16 representational dissimilarity matrix (RDM). These analysis steps were repeated per ROI. The separately estimated RDMs for each of the two-day scans were averaged across days. Third, we constructed a model RDM that contained the pairwise Euclidean distances

between entrepreneurs on the 2-D social hierarchy space (**Fig. 2A**). Last, we compared the neural RDMs estimated in each of the ROIs to our model RDMs using a rank correlation using Kendall's τ_A (Kendall, 1938). Statistical significance was calculated by constructing a null distribution by shuffling our data 1000 times at the subject level. Each subject's τ_A was compared to the mean of this permutation distribution and submitted to a one-sample Wilcoxon signed-rank test across participants. We reported the results with the family-wise error rate (FWE) to correct multiple comparisons by the number of ROIs using the Holm-Bonferroni method (Holm, 1979).

To test more rigorously the effect of Euclidean distance, which can be factorized into a 1-D distance of rank differences in the competence dimension and another 1-D distance of rank differences in the popularity dimension, we estimated the neural representation of the relationship between entrepreneurs in each of the two dimensions separately. We estimated the rank correlation (Kendall's τ_A) between the neural RDM which included neural activity patterns estimated from the same ROIs and the model RDMs which included the rank differences between entrepreneurs in each dimension. We further compared those rank correlations (τ_A) that represent to what extent a ROI contained the representation of 16 entrepreneurs' relationship in one of two social hierarchy dimensions, to test whether the neural representation of the relative ranks among entrepreneurs in one dimension was not different from that of the other dimension with the non-parametric two-sided Wilcoxon signed-rank test.

Multidimensional scaling

To visualize the representation of the social hierarchy structure built in HC and EC, we performed multidimensional scaling (MDS) analyses using the 'midscale' function in Matlab. The MDS allowed us to find the best way to arrange 16 individuals spatially in a 2-D space so that the distances between individuals in a 2-D space reflect their similarities between neural activity patterns. The distances between two individuals in the social hierarchy were measured by Mahalanobis distances between the neural activity extracted from the same HC and EC ROIs, as we computed for our RSA. The dissimilarity matrix of each participant was symmetrized by averaging the upper and the lower triangle, normalized by 0 to 1, averaged across hemispheres in the HC and EC ROI separately, and averaged across participants. We used the first two eigenvalues to reconstruct the representation in a 2-D space. Last, we tested to what extent the pairwise Euclidean distance between individuals in the MDS graph correlated with the true pairwise distance in the 4x4 social hierarchy structure using Pearson correlation.

In MDS, the locations of less sampled entities are more likely biased by noise compared to those of the entities sampled with higher frequency. Therefore, as can be seen below, the measures for face 1 and 16 in particular are expected to be less reliable than others, which may explain their relatively outlying estimated positions. We deliberately downsampled these faces to mitigate against subjects relying on heuristics during inferences, since Face 1 always lost on both dimensions and Face 16 always won on both dimensions. Importantly, we apply MDS only for visualization of the social hierarchy representation in the HC (**Fig. 2E**) and EC (**Fig. S1E**), not for statistical inferences, where we instead rely on rank correlations of RSA. We only report the Pearson correlation coefficients, along with the above caveat about sampling frequencies. The numbers of presentations of each of 16 faces during a single scan (three blocks of one day) were shown in brackets as follows:

Face 13 (42) Face 14 (21) Face 15 (16) Face 16 (8)

Face 9 (23)	Face 10 (58)	Face 11 (58)	Face 12 (19)
Face 5 (69)	Face 6 (45)	Face 7 (57)	Face 8 (18)
Face 1 (12)	Face 2 (69)	Face 3 (18)	Face 4 (43)

fMRI whole-brain analyses

In previous findings of rodent electrophysiology (Doeller et al., 2010; Fyhn, Molden, Witter, Moser, & Moser, 2004; Hafting et al., 2005), grid cells in the entorhinal cortex (EC) showed greater activity for navigating in the direction aligned with the mean grid orientation (ϕ) across grid cells compared to them for navigating in the direction misaligned with ϕ , which generated a specific pattern of six-fold symmetry (**Fig. 1B**). Recent studies have shown that a hexagonally symmetric grid-like code can be measured in the blood oxygen-level-dependent (BOLD) signal with fMRI in humans (Bao et al., 2019; Constantinescu et al., 2016; Doeller et al., 2010; Julian, Keinath, Frazzetta, & Epstein, 2018; Nau et al., 2018). We leveraged these established fMRI markers to test whether a grid-like code is used for inferred trajectories between entities in an abstract cognitive map to guide decision-making. Specifically, we searched for neural evidence for hexa-directional modulation of inferred decision trajectories through the reconstructed cognitive map of a 2-D social hierarchy that we hypothesized were used to compute the growth potential (GP) to guide decisions in the partner selection task: The $\overrightarrow{F0F1}$ vector at the time of F1 presentation and the $\overrightarrow{F0F2}$ vector at the time of F2 presentation (**Fig. 1E,F**).

To address this, we implemented a series of general linear models (GLM) to analyze fMRI data. First, we performed a whole-brain analysis to identify brain areas in which activity showed hexagonal symmetrical patterns according to the direction of the inferred trajectories (θ) (GLM1). Importantly, GLM1 was performed without any assumption of the subject-specific mean grid orientation (ϕ). Based on the results of GLM1, we define regions of interest (ROIs) in the brain in unbiased way for GLM2. Second, in GLM2, we tested the effects of hexadirectional modulations according to θ aligned with ϕ using a leave-one-and-out cross-validation (CV) procedure. Specifically, in GLM2, we estimated the putative grid orientation (ϕ) from one dataset (either two of 3 runs for within-day analyses or 1 of 2 days for between-day analyses) using an ROI defined independently based on GLM1 and applied the estimated ϕ to the other remaining dataset to test for the hexadirectional grid-like representation across the whole brain. Third, with GLM3, we identified brain areas encoding the variables guiding the partner selection decision including the growth potential (GP) and $|GP1-GP2|$, which is independent from both θ and ϕ . In addition, with GLM3, we further examined the relationship between neural encodings of the decision variables (GP) and individual differences in the effects of hexadirectional modulation of the inferred trajectories while controlling the overall accuracy of decision making. Last, with GLM4, we tested for complementary evidence of greater GP effects for the aligned compared to misaligned trajectories, consistent with a role of hexadirectional modulations in EC for guiding map-based decision making.

Defining a region of interest in an unbiased manner (GLM1)

With GLM1 we identified brain areas where activity showed hexagonally symmetric patterns according to the direction of the inferred trajectories (θ) but did not depend on the grid orientation (ϕ). The purpose of GLM1 was to functionally define ROIs in a way that was

unbiased from the putative grid orientation (ϕ) for future tests (see GLM2). GLM1 contained separate onset regressors for the onsets of F0 and the partner stimuli presentations (combining F1 and F2 presentations). The BOLD responses of those regressors were modeled with the 2.5s boxcar function. The regressors of F1 and F2 presentations were modulated with two parametric regressors which include the sine and cosine of the direction of inferred trajectory (θ) with a six-fold periodicity, that is, $\sin(6\theta)$ and $\cos(6\theta)$. These regressors produced parameter estimates, β_{\sin} and β_{\cos} with amplitude ($\sqrt{\beta_{\sin}^2 + \beta_{\cos}^2}$) for brain regions that are sensitive to hexagonal symmetry. When a partner face stimulus (Fx) followed F0, the direction of inferred trajectory θ_{F0Fx} was defined as follows:

$$\theta_{F0Fx} = \arctan\left(\frac{C_{Fx} - C_{F0}}{P_{Fx} - P_{F0}}\right)$$

where C_{Fx} and P_{Fx} indicate the level of competence and popularity dimensions of the individual, Fx respectively. We included the onset of button presses with a stick function as additional regressors of no interest to account for the potential motor related effects. The 6 motion parameters obtained during realignment were also entered into GLM1 as regressors of no interest. The orthogonalize function was turned off. All these regressors were convolved with the canonical hemodynamic response function (HRF). Using an F-test, we identified the brain areas in which BOLD signals were significantly modulated by a weighted sum of two regressors, $\beta_{\sin} \times \sin(6\theta) + \beta_{\cos} \times \cos(6\theta)$. This method we adopted here was established in a previous study (Constantinescu et al., 2016). All these regressors in these GLMs were convolved with the canonical hemodynamic response function (HRF). Individual whole brain contrast maps were z-scored and inputted to a one-sample t-test to test for group level effects using the z-statistic based on an asymptotic approximation (Jenkinson & Woolrich, 2004).

Considering that the F-statistics could be overestimated if the variance of the first level statistic was under-estimated, we only used GLM1 to create ROIs in an unbiased manner, but not for statistical inference. We defined ROIs for future tests in the entorhinal cortex (EC) (peak MNI coordinates at [x,y,z]=[26,-2,-36], z=2.62) and medial prefrontal cortex (mPFC) (peak MNI coordinates at [x,y,z]=[2,66,-4], z=5.58) (right panel in **Fig. S2A**). We defined the ROIs in the EC and mPFC based on the previous finding showing grid cells in these areas during virtual navigation in humans using pre-surgical single unit recordings (Jacobs et al., 2013) and previous grid-like fMRI effects in these regions (Constantinescu et al., 2016; Doeller et al., 2010). Based on the F-test results, we defined ROIs in the EC and mPFC within their anatomical masks defined in EC (Amunts et al., 2005; Zilles & Amunts, 2010) and mPFC (Neubert et al., 2015), while including the voxels showing the effects of six-fold symmetry at the threshold $z > 2.3$ which corresponds to $p = 0.01$.

Testing hexadirectional modulation aligned to the grid orientation (GLM2)

We performed the GLM2 to identify brain areas whose activity was modulated by the direction of inferred trajectories (θ) with six-fold symmetry aligned with the grid orientation (ϕ). To disentangle the effects of grid orientation (ϕ) from the direction of inferred trajectories (θ), we used a cross-validation (CV) procedure. This allowed us to identify hexadirectional modulations with consistent grid orientation (ϕ) across sessions.

The CV was performed as follows. We estimated ϕ from a separate dataset of each participant (called estimate set) and applied the participant-specific parameter ϕ to the other dataset (called test set) to test the effects of θ . For instance, the effect of θ in block 1 was

tested while applying ϕ estimated from blocks 2 and 3 acquired on the same day scan. Likewise, the ϕ for block 2 was estimated from blocks 1 and 3; and the ϕ for block 3 was estimated from blocks 1 and 2 (bottom left panel in **Fig. S2C**). The CV was possible because whereas the direction of inferred trajectories θ was consistent for the same trajectory across participants but varied across trajectories, the grid orientation ϕ was relatively consistent within a participant but varies across participants (see the top panel in **Fig. S2C**). We estimated the mean grid orientation ϕ in the functionally defined ROIs in the EC defined based on GLM1 (**Fig. S2B**), by adopting the same method proposed in previous studies (Constantinescu et al., 2016; Doeller et al., 2010; Nau et al., 2018). Specifically, we estimated ϕ in each voxel of the ROIs using the β coefficients for the sine and cosine regressors (β_{sin} and β_{cos}) in GLM1 as follows:

$$\phi = \frac{\arctan\left(\frac{\beta_{sin}}{\beta_{cos}}\right)}{6}$$

The mean grid orientation (ϕ) across each voxel in each of the ROIs was estimated using the circular mean (Berens, 2009), which was divided by six to compute ϕ in a range between $-\frac{\pi}{6}$ to $\frac{\pi}{6}$, thus taking the six-fold periodicity into account.

In GLM2, we modeled the BOLD signals of onsets of the partner stimuli (combining F1 and F2 presentations) with a 2.5 boxcar function. These events were modulated by a parametric regressor, $\cos(6[\theta - \phi])$, where θ was the hypothesized direction of inferred trajectories between individuals in each pair and ϕ was the mean grid orientation estimated from the separate estimate dataset. Additionally, the onsets of F0 and button presses were modeled as separate regressors with a stick function. The 6 motion regressors were also entered as regressors of no-interest. The orthogonalize function was turned off. All regressors in these GLMs were convolved with the canonical HRF. Last, the group-level effects were tested with a one-sample t-test.

In addition to the grid orientation consistency across sessions acquired within the same day, we further tested for the grid orientation consistency across sessions acquired on different days. Again, to test the effect of hexadirectional modulations in a statistically unbiased manner, we used a CV procedure. Specifically, the effect of the was tested on the fMRI sessions acquired in one of two days scans while applying the grid orientation, ϕ estimated in the EC activity from a session acquired in the other day scan (bottom right panel in **Fig. S2C**). This additional analysis allowed us to test whether the brain activity had the hexadirectional modulations in one day consistently showed the same effects to indicate the inferred trajectories in a different day. That is, we tested for a stable representation of the learned social hierarchy in the brain more than a week later.

Identifying neural encoding of decision-related variables (GLM3)

To make accurate decisions in each trial of the partner selection task, participants were required to compute GP1 for the F1 presentation and GP2 for the F2 presentation. Specifically, the GP of each pair corresponds to the size of the rectangle that is drawn by the rank positions of two entrepreneurs in the abstract 2-D social hierarchy space (as **Fig. 1F** shows). Subsequently, the decision should be guided by the difference between the two GPs. To capture this value comparison term, we use the absolute difference in GPs ($|GP1 - GP2|$).

We performed the GLM3 to examine the neural encodings of decision variables (GP and $|GP1 - GP2|$) which guide the partner selection decision. We modeled the BOLD signals

of onsets of the partner stimuli (combining F1 and F2 presentations) with a 2.5 boxcar function which were subsequently modulated by a parametric regressor, GP. We also modeled the BOLD signals at the time of decision-making with a stick function which was modulated by a parametric regressor, |GP1-GP2|. Additionally, the onset of F0 and the 6 motion regressors were also modeled with separate regressors in the same way of the GLM1. The orthogonalize function was turned off. All regressors in these GLMs were convolved with the canonical HRF.

The GP was dependent on the relative ranks of two entrepreneurs in both dimensions, but independent from the directions of the inferred trajectories (**Fig. S4A**, mean $r \pm SE = 0.04 \pm 0.03$). To rigorously test whether these two regressors explained independent variance in neural activity, as a control analysis, we performed an additional GLM in which the BOLD signals at F1 and F2 presentations were modulated by two parametric regressors, $\cos(6[\theta - \phi])$ and GP, in the same model. The effects of each of those two regressors were further compared with the effects of $\cos(6[\theta - \phi])$ in GLM2 and the GP effects in GLM3 to examine whether the effects of two regressors remained while controlling for the effect of the other regressor.

Interaction effects between alignments to the grid orientation and GP computations (GLM4)

In GLM 4 we tested the hypothesis that the GP effects depend on the alignments of the inferred trajectories to the EC grid orientation (ϕ) by comparing the GP effects between aligned trajectories and misaligned trajectories. Note that there was not a significant difference between GP regressor values for aligned and misaligned trajectories (**Fig. S4B,C**). In GLM4, we modeled the BOLD signals at the onset of the partner stimuli (combining both F1 and F2 presentations) separately for the trajectories aligned with the grid orientation (ϕ) from those misaligned trajectories with a 2.5s boxcar function. Both regressors were parametrically modulated by GP of the pairs. Consistently, we applied the same putative grid orientation, ϕ estimated in the EC ROI using a CV across sessions acquired within the same day, as before. Additionally, the onsets of F0 presentation and button presses were modeled with separate regressors with a stick function. The 6 motion parameters obtained during realignment were entered into the model as regressors of no interest. The orthogonalize function was turned off. All these regressors were convolved with the canonical HRF. We performed a contrast analysis between the GP effects in the aligned trajectories and those in the misaligned trajectories. The individual contrast maps were entered into a one-sample t-test to test group-level effects.

Group level inference

For all analyses using a GLM (GLM1,2,3, and 4), group-level effects were tested with a one-sample t-test. First, we performed group-level inference on *a priori* task-independent EC ROIs which were defined anatomically (Amunts et al., 2005; Zilles & Amunts, 2010). Second, we performed exploratory whole-brain analyses. To correct for multiple comparisons, we used permutation-based threshold-free cluster enhancement (TFCE) (Smith & Nichols, 2009). Specifically, the EC effects were reported at the threshold $p_{TFCE} < 0.05$ corrected within the independently defined anatomical ROI. For the whole-brain analyses, we reported results at the threshold $p_{TFCE} < 0.05$ corrected across the whole brain.

Region of interest analyses

Specificity for six-fold symmetry in ROIs

We conducted the regions of interest (ROI) analyses to test the hexadirectional patterns in alignment with the EC grid orientation (ϕ) within unbiased ROIs. Moreover, we visualized the sinusoidal modulations (**Fig. 1C**) in the activity in these ROIs as a function of the direction of inferred trajectories (θ) with six-fold periodicity. In addition, we performed control analyses to test whether the periodic effects in the ROIs were specific to six-fold but not to other control periodicities (four-, five-, seven- and eight-fold).

First, we tested the hexadirectional modulations on *a priori* ROIs in the EC and mPFC which we defined based on the GLM1 (**Fig. S2B**) independently from the grid orientation (ϕ). We tested this effect also on the anatomically defined EC and mPFC ROIs as well ((Amunts et al., 2005; Zilles & Amunts, 2010) for EC; (Neubert et al., 2015) for mPFC) (**Fig. S3A**) to confirm the effects did not depend on how we defined the ROIs. Second, because we used faces as stimuli in our task, we tested the hexadirectional modulations in the ROIs defined in the fusiform face areas (FFA). Note that the FFA was specifically engaged with the current task in response to the face stimuli. To define FFA ROIs independently from ϕ , we included the brain area activated in response to the face stimuli compared to fixations. The inclusive mask was defined at the threshold, $t_{20} > 3.6$ which corresponds to $p < 0.001$, uncorrected (peak MNI coordinate $[x, y, z] = [16, -34, -6]$, $t = 10.94$ for right FFA and $[x, y, z] = [-18, -30, -10]$, $t = 9.81$ for left FFA; **Fig. S3B**). Third, we defined ROIs in the brain areas where activity quantitatively encoded the GP, which included the bilateral TPJ and mPFC based on the GLM3 (**Fig. 4A; Table S3A**). The inclusive mask was defined at the threshold, $t_{20} > 3.6$ which corresponds to $p < 0.001$, uncorrected. With the TPJ and mPFC ROIs, we tested whether the neural encodings of GP (β GP) was modulated by the direction of the inferred trajectories (θ) with the six-fold periodicity aligned to the EC grid orientation (ϕ). Note that the functional definitions of the above ROIs were all independent from the subsequent statistical tests performed for inference.

To estimate the trajectory-by-trajectory neural activity, we modeled the BOLD signals at the time of onsets of every presentation of F1 and F2 stimuli as separate regressors with a 2.5s boxcar function. These event regressors were also modulated by a parametric regressor for the GP. The onsets of button presses were modeled with a separate regressor with a stick function. The 6 motion parameters were added as regressors of no-interests. All regressors were convolved with the HRF.

First, we extracted the BOLD signals activated at the time of every presentations of a partner stimulus (including both F1 and F2) in the EC, mPFC, and FFA ROIs. Second, we extracted the series of β signals encoding GP at the time of every partner stimulus presentation in the bilateral TPJ and mPFC ROIs. The series of the extracted BOLD signals and β GP were separately z-scored per ROI within each block. Based on the subject-specific EC grid orientation (ϕ) estimated with the CV across sessions acquired within a day, we categorized the z-scores into 12 equal bins of 30° according to the direction of the inferred trajectory (θ). The z-scores of each participant were then averaged within each of 12 bins \pm SE. This procedure produced six mean activity of the aligned trajectories and six mean activity of the misaligned trajectories per ROI.

In each of the ROIs, we performed nonparametric one-sided t tests to test the effects of hexadirectional modulations which were measured as differences in z-scores (the aligned – misaligned trajectories' bins). In addition, to test for the specificity of the six-fold periodicity, we performed the additional control analyses. Specifically, we performed two-sided paired t tests to compare the differences in z-scores computed using the model of the six-fold periodicity (effects of hexadirectional modulations) to the differences in z-scores computed

using the other model assuming the different level of periodicity including four-, five-, seven- and eight-fold (effects of modulations in other alternative periodicities). Last, we performed leave-one-bin-out (LOBO) analyses which allowed us to confirm that the effect of hexadirectional modulations in β GP was not driven by the activity of any specific bin. Specifically, with the LOBO analysis, we performed one-sided t-tests to examine whether the effects of hexadirectional modulations (difference in mean z-scores between the aligned and misaligned trajectories) were still greater than zero when the effects were computed from only 11 bins while excluding each of 12 bins (**Fig. S4F**).

Between-subject relationship between hexadirectional modulations in EC and the GP effects

To test for a relationship between hexadirectional modulation and GP effects, we addressed the following question: Do those participants who had a greater hexadirectional modulation in EC also show greater neural encoding of GP? The underlying rationale for this hypothesis is as follows. If the grid-like codes in the EC are important for accurate representations of inferred trajectories, it might further serve for accurate computations of the decision values in decision-making guided by a cognitive map. To test this hypothesis, we defined the subject-specific “gridness” score as to what extent the EC signal was modulated by the direction of inferred trajectories in six-fold periodicity. Therefore, the gridness score equals to the $\beta \cos(6[\theta - \phi])$ in EC, and then we tested whether individual differences in neural GP effects were explained by gridness scores. Specifically, using GLM3, we entered individual contrast maps of GP into a one-sample t-test to test the group-level effects and also entered the gridness score in EC as a (demeaned) covariate for the second-level regression analysis across participants. In addition to the gridness score, we also include the overall accuracy normalized in a range of 0 to 1 as an additional covariate to control for the choice accuracy. To test the between-subject effect in the brain areas engaged in the tasks requiring perspective taking, we tested this effect within the anatomically defined independent ROIs including bilateral TPJ (Mars et al., 2012) and mPFC (Neubert et al., 2015) (**Fig. 4C**).

The GP effects modulated by the inferred trajectories aligned to the EC grid orientation

We examined the effects of grid alignments in EC for making inferences of direct trajectories between individuals in the abstract social space on the activity in the brain areas encoding GP. We defined the ROIs in the mPFC and bilateral TPJ in which their activity correlates with GP while presenting F1 or F2 stimuli (GLM3) within the inclusive mask defined at the threshold $t_{20} > 3.6$ which corresponds to $p < 0.001$, uncorrected (**Fig. 4A**; **Table S3A**). In the ROIs, we extracted the β GP of the aligned trajectories and the β GP of misaligned trajectories which indicated to what extent the BOLD signals encoded the GP when the inferred trajectories were aligned and that for the trajectories misaligned to the EC grid orientation (ϕ). Note that the GP was independently defined from the direction of inferred trajectories (θ) regardless of the EC grid orientation (ϕ) (**Fig. S4A,B,C**, mean $r \pm SE = 0.04 \pm 0.03$). The GP, therefore, also independent whether the inferred trajectory was aligned or misaligned which obviates the concerns of the ROI selection biases. For each ROI, we performed a one-sample t test to confirm their activity encode GP (β GP > 0) but also a paired t-test to test whether the β GP was greater for the trajectories aligned to the EC grid orientation compared to the β GP of misaligned trajectories (**Fig. 4E**).