Do grid codes afford generalization and flexible decision-making?

Scientific question Grid-like codes in the human entorhinal cortex (EC) have been proposed to take part in learning and representing structural knowledge, which enables efficient planning and model-based decision-making. However, it has not yet been established whether these grid-like codes serve to transfer learning and to generalize previous experiences for inferences in novel situations. Do grid codes provide consistent context-invariant representations of structural relationships or do they flexibly switch their reference frame according to the current task demand to plan/discover optimal decision policy? Here we aim to reveal the constraints, conditions, and potential transformations of grid codes in the human brain that would afford flexible generalization.

Background The place cells in the hippocampus (HC) and grid cells in EC provide important neuronal codes for spatial navigation. The firing of grid cells tessellates space in regular hexagonal receptive fields (*1*, *2*). The grid cells may allow an agent to infer the structure (e.g. statistical relationships between positions, states, or entities) in an environment and decide whether to reuse or create a new cognitive map of the structure in a novel context. In two different but structurally similar environments, grid cells realign/reorient their codes in concert, which may allow the agent to generalize spatial codes between environments during navigation (*3*). Recent studies have suggested that such codes may reflect a general mechanism for organizing relationships between concepts or entities in abstract knowledge dimensions (*4*–*7*) and afford generalization of previous experiences in a novel situation (*8*–*10*). Recent computational models have proposed the idea of factorized representations in which the brain abstracts the latent task structure generalized across multiple contexts divorced from features specific to a single context (*11*–*14*). However, no experimental work has shown whether and how the grid codes represent the factorized low dimensional structural relationship and use this representation for generalizing previous experiences in a new context. In addition to EC, grid cells have also been found using pre-surgical electrocorticogram recordings in the medial prefrontal cortex (mPFC) and posterior cingulate cortex (PCC) in humans (*15*). Grid codes in these different areas may also serve different purposes: mPFC has greater connectivity to brain areas encoding task structure (*16*–*21*), while PCC is more connected with areas encoding action plans (*22*–*25*). However, little is known about whether the grid codes in these areas share the same coordinate system when the same relational codes guide different aspects of action selection or decision-making.

Challenge or controversy While theoretical models have proposed various mechanisms through which a grid code can arise in the brain (*26*–*30*), it remains elusive whether and how the grid codes support planning actions in a novel environment. Previous work on the subject has not resolved this issue because the agent learns the relational codes between entities and applies the knowledge to make decisions in the same environment, and thus does not need to develop a new decision policy. Instead, they could reuse the same decision policy not only for learning the environment but also for making novel inferences. Furthermore, though one team in our collaboration has identified the grid code in an abstract knowledge space of social hierarchy that was acquired by participants (*6*), evidence from the other team suggests that such a code may be absent in a conceptual space consisting of decisions between one fixed option and a varying second option (*31*). Though the two teams agree that grid codes might play a critical role in generalization in novel situations, we disagree on how they do so. On the one hand, they could do so through representation and abstraction of structural relationships over previous experiences (ultimately dependent on the previous learning policy) such that the same grid-codes are employed to transfer the learned structure to novel instances. On the other hand, they could be used for planning actions or making optimal decisions in the current task demands. This application allows for inferences of not-previously-experienced routes and discovering a new decision policy. If the first perspective is right, we would expect to find that training on trajectories that do not enable complete exploration of the transition structure of the space would abolish the grid representation. However, it would not if the second hypothesis is correct. Therefore, there are two main challenges to understanding the grid code's utility in cognition: a) determining whether grid codes use a context-invariant or context-dependent representation for discovering a decision policy for knowledge generalization and if the grid codes across the whole brain are bound to the same representation, and b) determining whether grid codes will be abolished in a situation where the transition structure is not directly learned through experience.

Competing hypotheses and proposed approach for resolution Both of the challenges arise out of two competing views on the following question: to afford generalization, how do the grid codes leverage a previously learned structure to find the solution for a novel problem? The first hypothesis is that the grid code is a low-dimensional representation of cached previous trajectories such as eigenvectors of the successor representation (SR), which is dependent on the policy that it is learnt over. The grid code would enable automatic generalization to new, similar contexts by reorienting/realigning, and transferring the structure of the previous environment to the new one, but any inferences of routes not learnt through previous policy would depend on other hippocampal mechanisms, such as simulation (*26*, *30*). According to this "learning" hypothesis, the grid codes should be explained by the allocentric representation of the task structure. For example, when a skier navigates the terrain, the grid code for her position will be based on a bird's eye view born out of her previous experiences navigating this area (such that the first-aid cabin is to the northeast of the lodge). The same EC grid-like codes should be used to indicate the vector between two points of the internal cognitive map irrespective of how this vector guides further action selection or decision-making. Moreover, activity of other brain areas showing the grid code would be aligned to the same EC grid orientation, suggesting the brain shares a single cognitive map. Because this code is determined by the policy during the initial learning experiences, we expect that when training experiences are limited and do not allow for experience of the complete transition structure of the task, the grid representations would be abolished.

The second alternative hypothesis is that the grid codes might take part in planning future actions and optimizing a decision policy in the current decision-making context. To make flexible decisions, the brain should switch between decision policies that determine how to map a state to an action. Under a different decision policy, the relationship between the same entities is mapping to different action values. Thus, for our skier, the grid code will depend on her policy based on her position (e.g., if she wants to get to the first-aid cabin from the lodge, then she must turn right, but for the other way around, she must turn left, and her grid code might be aligned to different reference frames according to the desired destination). A recent study shows that grid code is modulated by the current head direction of self-motion (*32*). To use a cognitive map to guide vector navigation in a novel situation, the grid codes should access the frontoparietal network in the brain to afford flexible decision-making and executive control (*33–35*). If this "planning" hypothesis is supported, we expect to find that the EC grid orientation estimated during decision-making in one context does not need to be aligned to that estimated in another context of decision-making, but systematically reoriented/realigned to the vector for efficient action selection or optimal decision-making. By extension from the second hypothesis, we also aim to test whether grid codes in either EC or outside of EC (or both) might play a role in translating the cognitive map into the low dimensional action values, complementing the recent findings of egocentric spatial representations in the posterior parietal cortex or secondary motor cortex (*22–25*). This hypothesis holds that grid codes should dynamically switch between decision policies, and that grid codes would exist even after incomplete training of transition structures because it would be used for planning and inference of new trajectories.

To test these competing hypotheses, we propose two experiments in which participants would be asked to generalize the previously learned structure into a novel situation and make different action selections or different decisions under a new decision policy. This approach is distinguished from previous work by decoupling the trajectories and the reference frames during initial learning experiences for building a cognitive map, and those for future planning and decision-making using this cognitive map**.** The first experiment dissociates spatial position from the task structure, to ask whether grid codes align to spatial position or goal demands. The second experiment instead puts attribute space positional coding into conflict with the frame of goal-related valuation.

In the first experiment, we test participants on a predictive inference task that requires learning an abstract task space that is dissociated from the physical representation of the space. In the first session, participants will learn a task where they spatially navigate a "shield" from a randomized start position to a location on a circular arena in which an "alien attack" is expected to occur (**Fig 1A**). The color of the shield is predictive of the attack location, with each of five colors predicting a different location on the arena, which the participants learn through experience. After each color-location association has been repeated a number of times, the attack locations will be rotated by 90 degrees with respect to the circular arena (**Fig 1B**). For the remainder of the task, the color-space associations will alternate between the first and second orientations. Notably, this manipulation changes the stimulus-response (location) contingency while preserving the relative spatial relationships between the attack locations for the different colors, which could in principle be able to be transferred between orientations. We test how well participants can use the abstract relationships they have learned to generalize to new applications by adding in novel colors to which they will have to apply the 90 degree rotation (in the second session), and by introducing novel degrees of rotations that the participants will have to apply to the originally learned colors (in the third session). To achieve better-than-expected accuracy on these novel applications before seeing any feedback is known as "zero-shot" learning. Critically, the randomized starting location in the experiment ensures that participants will experience a representative sampling of state transitions in the task and enables a complete SR of the task structure. We will test the necessity of representative sampling to the grid code by conducting an alternate version of the training task where the shield always starts at the same location (i.e., the center of the circle) in a different set of participants, while keeping all of the other aspects of the task the same.

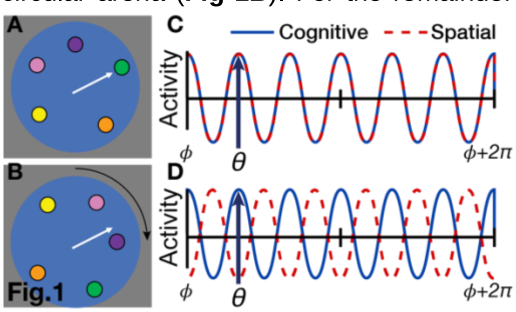

Using well-established previous methods to look at brain regions showing hexagonal symmetry in fMRI (*4*, *6*, *36*), the 90-degree rotation allows us to perfectly dissociate grid-like fMRI representations observed during task performance into spatial and abstract categories. The cognitive representation of the task space is aligned relative to the color locations and thus is re-oriented between blocks, whereas the spatial representation remains the same (**Fig 1C and 1D**). If the grid code represents the cognitive task space, the "learning" hypothesis holds that it would realign and automatically transfer the structural information from the original context (no rotation) to the new context (e.g., the 90-degree rotation), and its phase consistency should be correlated with the success on zero-shot learning. Also, various regions that show the grid code (mPFC, EC, and PCC) would all be aligned to this cognitive representation of the task (as opposed to the spatial representation). In the same-starting-location version of the task, this account predicts that grids will not emerge, and that instead behavior will be guided by value signals over possible trajectories, analogous to the lack of grids observed in previous study by one of our teams (*31*). In contrast, the "planning" hypothesis would hold that the grid code is used for planning independently in each rotation, irrespective of the starting point manipulation. In this case, the structural information of the task should not automatically transfer between contexts and enable zero-shot learning. The grid code in EC, mPFC, and PCC may differ in whether they show alignment to the spatial or the cognitive representation, depending on whether they are involved in higher order planning (mPFC, EC), or motor planning (PCC). In the same-starting-location version of the task, the planning hypothesis would predict that a grid representation would still be intact, due to its role in planning that would allow inference of the relationships between the colors regardless of experience during training.

In the second experiment, we test whether and how the grid-codes serve to generalize the previously learned cognitive map to make a decision which requires one to develop a novel decision policy. During behavioral training (**Fig.2A**), participants are trained on an abstract 4x4 ‘wine space’ characterized by two orthogonal dimensions of four attributes: a taste dimension (sweet to dry) and a body dimension (light to full). During learning blocks of training, participants learn piecemeal the relative attributes of wines from feedback-based binary comparisons. Critically, participants only compare two wines that are different by one level in the task-relevant attribute, given as colored cues (**Fig.2A**). Importantly, they learn the relative state of a pair of wines along one attribute at a time in separate days (e.g. On Day 1, only the taste dimension is experienced; In Day 2, only the body dimension is experienced). They are never asked to combine two dimensions during training. Moreover, participants are not presented with the overall graph structure, nor are they explicitly encouraged to resolve this task spatially. Instead, they might be able to infer the structure with transitive inferences. During fMRI (Day 3, **Fig.2B**), participants are asked to choose a wine which is preferred in the markets of different countries, each of which prefers either one or two taste attributes in their wines. For example, Spain prefers dry wines while France prefers sweet and bold wines. Therefore, the participants should switch between different decision policies to map a wine (state) into a different value (action value) in different countries. Note that participants have only used the decision policies aligned to one dimension while learning the cognitive map (1D hierarchy; **Fig.2A**), and never have a chance to develop a policy aligned to the diagonal axis which is required for choosing a wine preferred in a country in which the market value is determined by the relative states in two dimensions (2D hierarchy; **Fig.2B**). As with the first experiment, we plan to identify the brain regions showing the hexadirectional symmetry according to the direction of inferred trajectories over the 4×4 abstract wine space (*4*, *6*, *36*) (**Fig.2C**). If the “learning” hypothesis is supported, the grid code should be used for representing a low dimensional projection of the SR, or task structure. Accordingly, there are two possible expected results. First, if participants only could build two separate SR per dimension, all decision trajectories should be aligned to a single decision axis with zero variance between angles in decision trajectories, which will make the grid codes abolished (**Fig.2D left**). Second, if participants could find an optimal way to assemble two independent dimensions and represent a 2D cognitive map which comprises two orthogonal dimensions, as we found in the previous studies using the similar paradigm (*6*, *21*), we expect to see the grid codes modulated by the direction of the inferred vectors between wines but consistent across contexts (countries) (**Fig.2D right**). If the “planning” hypothesis is supported, we expect not only to see the grid codes but changes of their basis (the orientation of the grid codes) realigned with the decision policy according to the taste preference of each country (**Fig.2E**). The changes in basis of the grid codes will be tested not only within each of regions of interest in EC, mPFC, and PCC, but also between regions (e.g. test whether the grid code in PCC is aligned to the grid orientation estimated from EC). To take it one step further, we will conduct representational similarity analysis (RSA) on fMRI data. If the grid codes across the whole brain share a single allocentric coordinate to provide a context invariant grid code, we expect to find that the representational dissimilarity matrix (RDM) will only be explained by the pairwise 2D distance (e.g. Euclidean) between wines in the cognitive map (**Fig.2D left**). On the contrary, if the basis of the grid codes is flexibly changed across countries realigned to the optimal decision policy in the current task, we expect to find that the changes in pattern dissimilarity explained by the function of value differences between wines (**Fig.2E**).

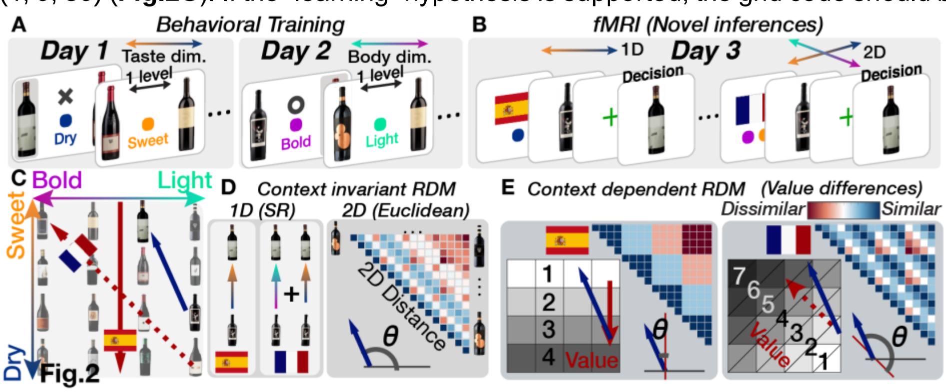

Concrete outcomes We aim to understand the computational mechanism underlying the ability to generalize previous experiences to novel problems. We proposed two experiments that were tailored to test the hypothetical roles of grid codes in the human brain. We will be able to answer the following questions: Do grid codes inform context-invariant (allocentric) relational codes across contexts or context-specific (egocentric) vectors for optimal decision-making? If the latter, is the new basis aligned to either the axis for selecting planned actions or to the gradient of the decision value distributions under a new policy, or both? Could the grid code be abolished with an incomplete transition structure? By providing experimental evidence, we aim to contribute to a better understanding of how grid codes are employed to format structural knowledge enable zero-shot inferences in novel situations.

Benefit to the community By providing solid experimental evidence of the role of grid codes in the generalization of abstract knowledge in the general domain, we believe that this work will have potential positive influences in integrating reinforcement learning and representational learning perspectives and more broadly in cognitive computational neuroscience society as a whole. Behavioral flexibility across multiple contexts is a hallmark of human intelligence yet challenging to machines (*37*, *38*) c.f. (*39*). Our work will contribute to designing intelligent systems that represent an abstract knowledge structure from experiences and use this structure to facilitate transfer learning.

References

1. M. Fyhn, S. Molden, M. P. Witter, E. I. Moser, M. B. Moser, Spatial representation in the entorhinal cortex. *Science*. **305**, 1258–1264 (2004).
2. T. Hafting, M. Fyhn, S. Molden, M.-B. B. Moser, E. I. Moser, Microstructure of a spatial map in the entorhinal cortex. *Nature*. **436**, 801–806 (2005).
3. M. Fyhn, T. Hafting, A. Treves, M. B. Moser, E. I. Moser, Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* (2007), doi:10.1038/nature05601.
4. A. O. Constantinescu, J. X. O'Reilly, T. E. J. Behrens, Organizing conceptual knowledge in humans with a gridlike code. *Science*. **352**, 1464–1468 (2016).
5. X. Bao, E. Gjorgieva, L. K. Shanahan, J. D. Howard, T. Kahnt, J. A. Gottfried, Grid-like Neural Representations Support Olfactory Navigation of a Two-Dimensional Odor Space. *Neuron*. **102**, 1066-1075.e5 (2019).
6. S. A. Park, D. S. Miller, E. D. Boorman, Novel Inferences in a Multidimensional Social Network Use a Grid-like Code. *bioRxiv*, 1–13 (2020).
7. H. Eichenbaum, N. J. Cohen, Can We Reconcile the Declarative Memory and Spatial Navigation Views on Hippocampal Function? *Neuron*. **83**, 764–770 (2014).
8. T. E. J. Behrens, T. H. Muller, J. C. R. Whittington, S. Mark, A. B. Baram, K. L. Stachenfeld, Z. Kurth-Nelson, What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron*. **100**, 490–509 (2018).
9. J. L. S. Bellmund, P. Gärdenfors, E. I. Moser, C. F. Doeller, Navigating cognition: Spatial codes for human thinking. *Science*. **362** (2018), p. eaat6766.
10. J. Lisman, G. Buzsáki, H. Eichenbaum, L. Nadel, C. Ranganath, A. D. Redish, Viewpoints: how the hippocampus contributes to memory, navigation and cognition. *Nat. Neurosci.* **20**, 1 (2017).
11. A. Banino, C. Barry, B. Uria, C. Blundell, T. Lillicrap, P. Mirowski, A. Pritzel, M. J. Chadwick, T. Degris, J. Modayil, G. Wayne, H. Soyer, F. Viola, B. Zhang, R. Goroshin, N. Rabinowitz, R. Pascanu, C. Beattie, S. Petersen, A. Sadik, S. Gaffney, H. King, K. Kavukcuoglu, D. Hassabis, R. Hadsell, D. Kumaran, Vector-based navigation using grid-like representations in artificial agents. *Nature*. **557**, 429–433 (2018).
12. T. J. Madarasz, T. E. Behrens, Better transfer learning with inferred successor maps. *Adv. Neural Inf. Process. Syst. 32 (NIPS 2019)*, 9029–9040 (2019).
13. P. Piray, N. D. Daw, A common model explaining flexible decision making, grid fields and cognitive control. *bioRxiv* (2019), doi:10.1101/856849.
14. J. C. Whittington, T. H. Muller, S. Mark, G. Chen, C. Barry, N. Burgess, T. E. Behrens, The Tolman-Eichenbaum Machine: Unifying space and relational memory through generalisation in the hippocampal formation. *bioRxiv*, 770495 (2019).
15. J. Jacobs, C. T. Weidemann, J. F. Miller, A. Solway, J. F. Burke, X.-X. X. Wei, N. Suthana, M. R. Sperling, A. D. Sharan, I. Fried, M. J. Kahana, Direct recordings of grid-like neuronal activity in human spatial navigation. *Nat. Neurosci.* **16**, 1188–1190 (2013).
16. R. C. Wilson, Y. K. Takahashi, G. Schoenbaum, Y. Niv, Orbitofrontal cortex as a cognitive map of task space. *Neuron*. **81**, 267–279 (2014).
17. S. C. Y. Chan, Y. Niv, K. A. Norman, A Probability Distribution over Latent Causes, in the Orbitofrontal Cortex. *J. Neurosci.* **36**, 7817–7828 (2016).
18. N. W. Schuck, M. B. Cai, R. C. Wilson, Y. Niv, Human Orbitofrontal Cortex Represents a Cognitive Map of State Space. *Neuron*. **91**, 1402–1412 (2016).
19. A. M. Wikenheiser, Y. Marrero-Garcia, G. Schoenbaum, Suppression of Ventral Hippocampal Output Impairs Integrated Orbitofrontal Encoding of Task Structure. *Neuron*. **95**, 1197-1207.e3 (2017).
20. J. Zhou, M. Montesinos-Cartagena, A. M. Wikenheiser, M. P. H. Gardner, Y. Niv, G. Schoenbaum, Complementary Task Structure Representations in Hippocampus and Orbitofrontal Cortex during an Odor Sequence Task. *Curr. Biol.* **29**, 3402-3409.e3 (2019).
21. S. A. Park, D. S. Miller, H. Nili, C. Ranganath, E. D. Boorman, Map Making: Constructing, Combining, and Inferring on Abstract Cognitive Maps. *Neuron*, 810051 (2020).
22. C. Summerfield, F. Luyckx, H. Sheahan, Structure learning and the posterior parietal cortex. *Prog. Neurobiol.* **184** (2020), p. 101717.
23. J. M. Olson, J. K. Li, S. E. Montgomery, D. A. Nitz, Secondary Motor Cortex Transforms Spatial Information into Planned Action during Navigation. *Curr. Biol.* **30**, 1845-1854.e4 (2020).
24. A. A. Wilber, B. J. Clark, T. C. Forster, M. Tatsuno, B. L. McNaughton, Interaction of egocentric and world-centered reference frames in the rat posterior parietal cortex. *J. Neurosci.* **34**, 5431–5446 (2014).
25. J. R. Whitlock, G. Pfuhl, N. Dagslott, M. B. Moser, E. I. Moser, Functional Split between Parietal and Entorhinal Cortices in the Rat. *Neuron*. **73**, 789–802 (2012).
26. Y. Dordek, D. Soudry, R. Meir, D. Derdikman, Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *Elife*. **5** (2016), doi:10.7554/eLife.10094.
27. L. M. Giocomo, M. B. Moser, E. I. Moser, Computational models of grid cells. *Neuron*. **71**, 589–603 (2011).
28. M. Klukas, M. Lewis, I. Fiete, Efficient and flexible representation of higher-dimensional cognitive variables with grid cells. *PLoS Comput. Biol.* **16**, e1007796 (2020).

29. R. M. Mok, B. C. Love, A non-spatial account of place and grid cells based on clustering models of concept learning. *Nat. Commun.* **10** (2019), doi:10.1038/s41467-019-13760-8.
30. K. L. Stachenfeld, M. M. Botvinick, S. J. Gershman, The hippocampus as a predictive map. *Nat. Neurosci.* **20**, 1643–1653 (2017).
31. S. Lee, L. Q. Yu, C. Lerman, J. W. Kable, Subjective value, not a grid-like code, describes neural activity in ventromedial prefrontal cortex during value-based decision-making. *bioRxiv*, 759951 (2019).
32. K. Gerlei, J. Passlack, H. Stevens, I. Papastathopoulos, M. F. Nolan, Grid cells encode local head direction. *bioRxiv*, 681312 (2020).
33. M. Donoso, A. G. E. Collins, E. Koechlin, Foundations of human reasoning in the prefrontal cortex. *Science*. **344**, 1481–1486 (2014).
34. A. G. E. Collins, M. J. Frank, Cognitive control over learning: Creating, clustering, and generalizing task-set structure. *Psychol. Rev.* **120**, 190–229 (2013).
35. D. Badre, D. E. Nee, Frontal Cortex and the Hierarchical Control of Behavior. *Trends Cogn. Sci.* **22** (2018), pp. 170–188.
36. M. Nau, T. Navarro Schröder, J. L. S. Bellmund, C. F. Doeller, Hexadirectional coding of visual space in human entorhinal cortex. *Nat. Neurosci.* **21**, 188–190 (2018).
37. T. Kriete, D. C. Noelle, J. D. Cohen, R. C. O'Reilly, Indirection and symbol-like processing in the prefrontal cortex and basal ganglia. *Proc. Natl. Acad. Sci.* **110**, 16390–16395 (2013).
38. B. Lake, M. Baroni, in *35th International Conference on Machine Learning, ICML 2018* (International Machine Learning Society (IMLS), 2018), vol. 7, pp. 4487–4499.
39. F. Hill, A. Lampinen, R. Schneider, S. Clark, M. Botvinick, J. L. McClelland, A. Santoro, Environmental drivers of systematicity and generalization in a situated agent (2019).

Group roles

Learning, Memory and Decision Lab, Brown University

Linda Q. Yu: Linda is a postdoctoral research associate, and will be a junior group leader. She will lead the first experiment, as well as jointly lead the GAC workshop and position paper.

Matthew R. Nassar: Matt is an assistant professor and principal investigator of the lab. He will be a senior group leader and advise over the GAC process.

Learning and Decision-Making Lab, University of California, Davis

Seongmin A. Park: Seongmin is a project scientist. He will be a junior group leader. He will lead the second experiment, and jointly lead the GAC workshop and position paper.

Sarah C. Sweigart: Sarah is a Ph.D. student. She will lead the second experiment of this GAC process.

Erie D. Boorman: Erie is an assistant professor and principal investigator of the lab. He will be a senior group leader and advise over the GAC process.

Statement of Commitment The following members agree to the GAC process on the proposed topic, including: 1) incorporation of feedback to the proposal and potential addition of new members; 2) running an online workshop for CCN 2020; 3) writing the position paper; 4) attending and presenting progress at CCN 2021.

LQY	MRN	SAP	SCS	EDB
linda_yu@brown.edu	matthew_nassar@brown.edu	apark@ucdavis.edu	ssweigart@ucdavis.edu	edboorman@ucdavis.edu