Generalizable Relational Inference with Cognitive Maps in a Hippocampal Model and in Primates

Jaedong Hwang* **Sujaya Neupane*** **Mehrdad Jazayeri†** **Ila Fiete†**
Massachusetts Institute of Technology
{jdhwang, sneupane, mjaz, fiete}@mit.edu

Abstract

We investigate the role of cognitive maps and hippocampal-entorhinal architecture in a mental navigation (MNAV) task by conducting experiment in humans, monkeys and neural network models. Humans can generalize their mental navigation performance to untrained start-target landmark pairs in a given landmark sequence and also rapidly adapt to new sequences. The model uses a continuous-time recurrent neural network (CTRNN) for action decisions and a hippocampal-entorhinal model network, MESH (Memory network with Scaffold and Heteroassociation), for encoding and learning maps. The model is first trained on a navigation-to-sample (NTS) task and tested on MNAV task where no sensory feedback is available, across five different environments (i.e. landmark sequences). The CTRNN with MESH solves MNAV task by reconstructing the next image via path integration and vastly outperforms the model with CTRNN alone. In both NTS and MNAV tasks, MESH-CTRNN model shows better generalization to untrained pairs within each environment and faster adaptation to new environments. Like humans, monkeys also exhibit generalization to untrained landmark pairs in MNAV task. We compared the neural dynamics in monkeys' entorhinal cortex to the dynamics of CTRNN and found behaviorally relevant periodic signals in both. The study demonstrates the importance of hippocampal cognitive maps in enabling data-efficient and generalizable learning in the brain.

1 Introduction

Cognition involves organizing experiences into retrievable knowledge for novel mental computations, which is achieved through cognitive maps encoding spatial, temporal, and abstract relationships. Spatial contexts have been extensively studied, with sensory experiences driving spatially selective responses in the hippocampus and entorhinal cortex (EC) [5]. When animals encounter a new task that conceptually matches a previously seen task, it is common to observe rapid learning even if surface-level details and inputs differ. Such learning is believed to involve a transfer of conceptual understanding to the new task facilitated by the coding principle of the spatial cells reported in the hippocampal formation [1]. However, neural models of such generalization are lacking.

Recently, Neupane *et al.* [4] showed cognitive map being endogenously recruited in monkeys' entorhinal cortex to solve a mental navigation task. We extended this study by collecting human behavioral data on the mental navigation task and documenting three types of generalization as shown in Figure 1. We then built artificial neural network models that exhibited all three generalizations and showed neural dynamics similar to that in monkey EC. Conventional recurrent networks can perform the mental navigation task but fail to generalize. We hypothesize that a structured neocortical-entorhinal-hippocampal circuit, the Memory Scaffold with Heteroassociation (MESH) adapted

*: equal contribution, †: equal correspondence

Associative Memory & Hopfield Networks in 2023. NeurIPS 2023 workshop.

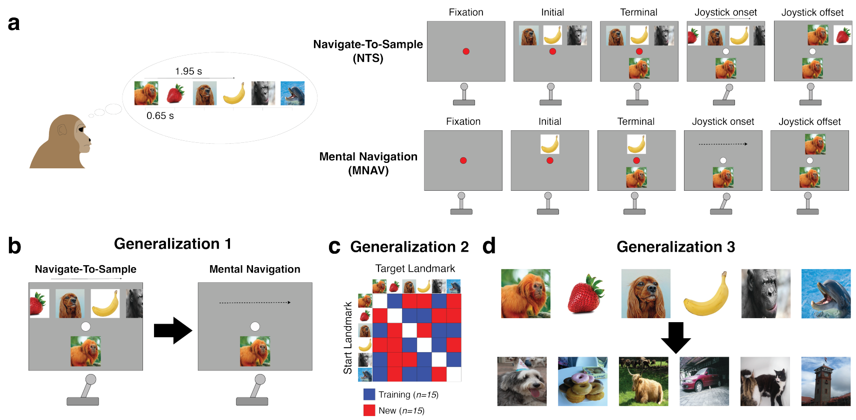

Figure 1: (a) Navigate-To-Sample (NTS) and Mental Navigation (MNAV) Tasks. An agent moves the joystick toward a target landmark presented on the bottom of the screen from a start landmark. In NTS, the current landmark is visible while it is invisible in MNAV. The distance between landmarks is fixed as 0.65s and the landmark sequence is also fixed in each environment. There are three types of generalization of this task; (b) from NTS to MNAV, (c) training start-target landmark pairs to new start-target pairs during MNAV, and (d) rapid adaptation to new landmark sequences after learning task.

from Sharma *et al.* [7] with grid cell modules can achieve the generalizations observed in human participants. We build a multi-region brain model, using a continuous-time recurrent policy neural network (CTRNN) to decide actions and the hippocampal MESH network to encode and learn maps. The learning rules in MESH are online and associative, based on velocity input and external cues. The outputs of MESH drive the action network. We sequentially trained the model on five landmark sequences in the navigate-to-sample setting. Our model achieved the same performance in visual and mental navigation tasks while CTRNN without MESH failed at mental navigation. The model exhibited better generalization to unseen pairs in each environment and adapted to new environments faster than the baseline. Furthermore, We found the periodicity patterns and distance-direction coding in the internal dynamics of the network that are also found in the neural recording of the entorhinal cortex [4]. Our work is thus a step toward a whole-system understanding of how the brain performs highly data-efficient and generalizable learning.

2 Mental Navigation Task

Neupane *et al.*[4] developed a mental navigation task for monkeys. Agents are trained on a sequence of six landmark images (Figure 1). Given a start and a target landmark, they must use a joystick to move between them (navigate-to-sample or NTS task). After reaching a performance criterion in NTS, the monkeys were introduced to the mental navigation task (MNAV). In MNAV, the image sequence was occluded and only the start and the target landmark were visible before joystick deflection. The sequence was hidden throughout the trial, including after the joystick deflection. To solve the task, the animals had to rely on their memory of relative landmark positions and navigate without sensory feedback. Monkeys successfully learned to perform the MNAV task, and the produced vectors closely matched the actual vectors in terms of magnitude and direction.

3 Behavioral Study

We extended the behavioral experiments to humans. Instead of a single six-landmark sequence as in Neupane *et al.*[4], we used three nine-landmark sequences. We randomly select 27 classes of images having distinctive objects from MSCOCO [3] dataset as landmarks. We collected data from 7 participants and they received a small gift regardless of their performance. Please refer to Appendix B for a detailed procedure. We found that human subjects can rapidly generalize to MNAV after training on NTS (Figure 2a) and to unseen start-target pairs that are not used during training (Figure 2b). Moreover, Figure 2c shows that the performance during the initial two blocks is improved for each new subsequent environment, suggesting that subjects learned the task structure and rapidly adapted to a new landmark sequences.

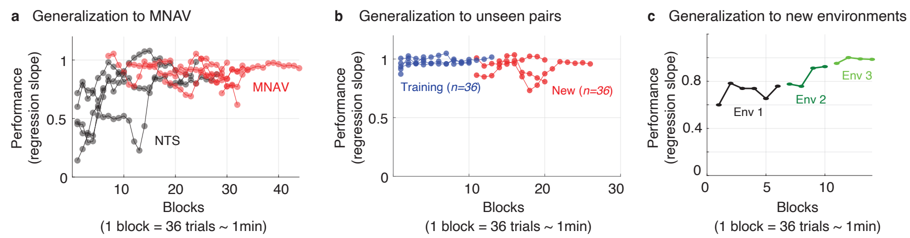

Figure 2: Results of behavioral experiments with human subjects. (a) NTS and MNAV performance in the first session. Humans can generalize from NTS to MNAV and memorize the landmark sequence the next day. (b) Humans can generalize from training pairs to training pairs. (c) NTS performance in each environment.

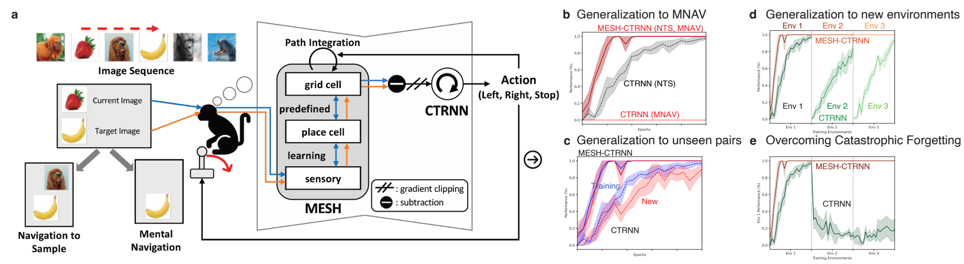

Figure 3: (a) The agent explores the sequence of images with blank intervals. On each trial, the input to the network is a pair of start and target images. The images are encoded as grid codes by MESH and fed into CTRNN to decide an action; move in one direction or stop the trial. As in the animal experiment, the start image is continuously updated in navigation-to-sample task but masked in the mental navigation task. (b) Success rate of NTS and MNAV in environment 1 of our MESH-CTRNN (solid line) and the baseline, CTRNN (dashed line). (c) Performance of two models about training and new pairs in environment 1 NTS setting. Performance of three environments in NTS setting (d) and the first environment while training subsequent environments (e). The lines and shades denote average and standard error, respectively.

4 Method

We built a multi-region brain model, using a CTRNN-based network to decide actions, and the entorhinal-hippocampal MESH network [7] to encode and learn maps that associate observations with the grid cell scaffold. Figure 3a shows the architecture of the model. The learning rules for map formation in MESH run online, based on velocity (action) input and external cues (sensory inputs). MESH is composed of three layers; sensory (input), place cell, and grid cell layers. The grid code (phase code) is formulated as a k-hot vector imposed by local recurrent inhibition, where k is the number of modules and each module has a different period. All codes are paired with place cell activation before training and the sensory input is associated with each grid code by pseudo-inverse learning between the input and place cells. Please refer to Appendix A and Sharma *et al.*[7] for more details. The association between the current image and the grid code is made via this learning mechanism during NTS. Given the grid codes of both current and target images, the CTRNN predicts the actions - move left, move right or stop.

In MNAV, MESH first associates both the start and the target images with corresponding grid codes. Upon taking an action, MESH infers the subsequent code for every subsequent image via path integration without having to refer to the landmark. The output of MESH drives the action network and the action is fed into the grid cell layer as a velocity input. Consequently, the proposed model can retrieve the correct grid representation for each image during the mental navigation task.

5 Model Experiments

5.1 Comparison with Vanilla Model

We sequentially trained the model in three different environments only in navigation-to-sample (NTS) using ground-truth actions. Each environment has six different landmarks with the same size of intervals similar to [4]. Among all possible pairs in each environment, 80% pairs are used in training

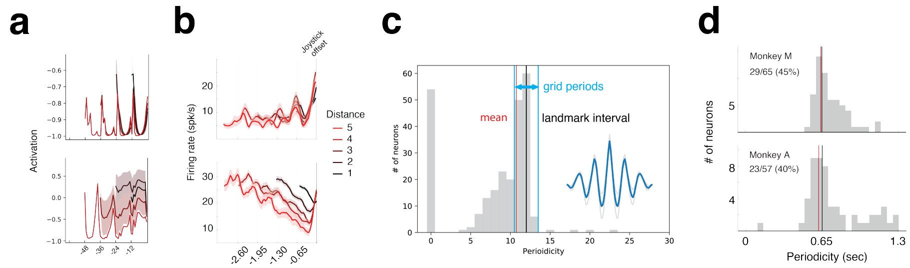

Figure 4: (a) activity changes of two neurons in CTRNN and (b) firing rate of two neurons aligned to joystick offset, color-coded by temporal distance. Histogram of periodicity of auto-correlation of each neuron in CTRNN (c) and entorhinal cortex. (b) and (d) are adapted from Figure 2 in Neupane *et al.* [4] There is a clear periodicity in neural recording (c) and auto-correlation.

(Training pairs) and the others are only used in testing (New pairs). The model was tested on both NTS and MNAV using all pairs. We employed the success rate to evaluate performance in all conditions: both navigation with Training and New start-target pairs. All experiments were conducted five times with different random seeds. Please refer to Appendix C for experimental details.

We demonstrate the success rates in the first environment for both the NTS and MNAV conditions in Figure 3b, and for Training and New pairs in the NTS condition in Figure 3c. Figure 3d illustrates the success rates in three different environments that were sequentially trained. Figure 3e shows the success rate in the first environment while training in other environments. Overall, MESH-CTRNN model can perform mental navigation with three distinct types of generalization observed in human subjects without catastrophic forgetting. In contrast, CTRNN without MESH can learn the task albeit slower but fails at generalizations and exhibits catastrophic forgetting. Such superior performance is because of the biologically inspired embedding space of grid phases in MESH architecture that can encode all observations efficiently.

5.2 Comparison between Model and Monkey Physiology

We compare the internal dynamics of the model and that of the entorhinal cortex of two monkeys performing mental navigation task (adapted from from Neupane *et al.* [4]). We analyzed the hidden states in CTRNN and found periodic patterns of activity in the hidden units 4a. The most prominent periodicity was the same as the image interval as shown in Figure 4c. We recently reported such a behaviorally relevant periodicity in monkey's entorhinal cortex as shown in Figure 4b and Figure 4d.

6 Discussion

We trained humans and a multi-brain region neural network model with a hippocampal-entorhinal scaffold network, MESH [7] to perform a novel mental navigation task [4]. Both showed three types of generalizations: (1) generalization to mental navigation by training on visual navigation (2) generalization to unseen start-target landmark pairs and (3) generalization in new environments with minimal learning. We also showed that the model's internal dynamics are similar to the neural recording in monkeys' entorhinal cortex which showed behaviourally relevant periodic signals.

We sequentially trained the model on five different environments and tested how quickly it could adapt to new environments and whether it could overcome catastrophic forgetting. The model could solve the purely mental version of NTS after training only on NTS, by using path integration to reconstruct the next grid state without referring to visual stimuli. It could also overcome catastrophic forgetting and learn new environments instantly (one-shot) by projecting observations into its structured embedding space based on the grid cell code.

One notable design feature of our model is the nature of grid code made available for the downstream CTRNN. We convert the k-hot grid phase into an integer code before feeding it as an input to CTRNN. The network performance was dependent on this choice among many other coding choices. One way to interpret this observation is to view the grid coding scheme as a decoding problem. As such, there must exist a brain area that reads out grid phase activation and transforms it into an integer code.

In sum, our model exhibits a high degree of alignment with behavior and neurophysiology and thus offers a rich test-bed to perform perturbation experiments and generate hypotheses.

Acknowledgments

SN is supported by NSERC PDF-516867-2018, FRQNT B3X-258512-2018. MJ is supported by NIH (NIMH-MH129046), Paul and Lilah Newton Brain Science Award, and the McGovern Institute. IF is supported by the Office of Naval Research, and the Howard Hughes Medical Institute (HHMI), and NIH (NIMH-MH129046).

References

[1] Timothy EJ Behrens, Timothy H Muller, James CR Whittington, Shirley Mark, Alon B Baram, Kimberly L Stachenfeld, and Zeb Kurth-Nelson. What is a cognitive map? organizing knowledge for flexible behavior. *Neuron*, 100(2):490–509, 2018.

[2] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2015.

[3] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.

[4] Sujaya Neupane, Ila R Fiete, and Mehrdad Jazayeri. Vector production via mental navigation in the entorhinal cortex. *bioRxiv*, pages 2022–12, 2022.

[5] John O'Keefe and Jonathan Dostrovsky. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain research*, 1971.

[6] L Personnaz, I Guyon, and G Dreyfus. Information storage and retrieval in spin-glass like neural networks. *Journal de Physique Lettres*, 46(8):359–365, 1985.

[7] Sugandha Sharma, Sarthak Chandra, and Ila Fiete. Content addressable memory without catastrophic forgetting by heteroassociation with a fixed scaffold. In *ICML*, 2022.

[8] Jonathan Tapson and André van Schaik. Learning the pseudoinverse solution to network weights. *Neural Networks*, 45:94–100, 2013.

Appendix

A MESH

Memory Scaffold with Heteroassociation (MESH) [7] is a content-addressable memory (CAM) model that stores vectors as fixed points of its dynamics and reconstructs them from noisy cues. It is composed of 'features', 'hidden states', and 'labels', that are originally designed for the memory cliff problem; if the number of stored vectors exceeds a threshold, the model fails to reconstruct all vectors. MESH constructs a fixed scaffold of pre-defined content-independent fixed points (*labels*), which are then used to store the content-laden patterns through hetero-associative learning.

The MESH architecture is similar to the hippocampus and the entorhinal cortex relation by corresponding hidden states and labels to place cell and grid cell layers, respectively. The place cell layer $\mathbf{p} \in \{-1, +1\}^{N_P}$ represents an N_P dimensional binary vector, and the grid cell layer $\mathbf{g} \in \{0, 1\}^{\sum \lambda_i}$ is defined as the concatenation of λ_i dimensional one-hot vectors, where λ_i is the period of i-th grid cell.

Before initiating the experiments, the memory scaffold, including the states of grid and place cells and their interconnections, is pre-defined. The matrix that projects the grid cell layer to the place cell layer, denoted as $\mathbf{W}_{PG}$, is set up in a random manner ensuring that it retains a one-to-one projection. Conversely, the matrix leading from the place cell layer back to the grid cell layer is adapted through Hebbian learning. This ensures that an active place cell, which defines a specific place code, is linked to simultaneously active grid cells, which represent an associated grid code:

$$\mathbf{W}_{GP} = \frac{1}{|\mathbf{N}|} \sum_{\mu=1}^{\mu=N} \mathbf{g} \cdot (\mathsf{sign}(\mathbf{W}_{PG} \cdot \mathbf{g}))^T, \tag{1}$$

where N is the number of training patterns (vectors).

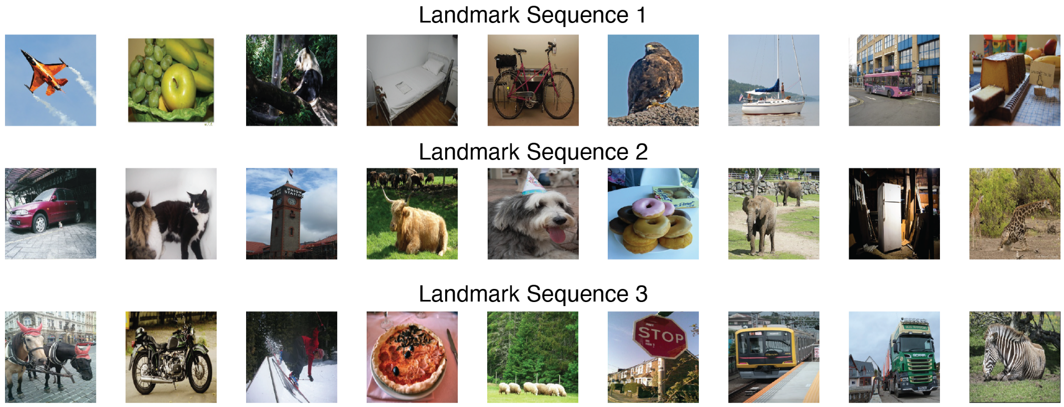

Figure 5: Landmark sequences. There are nine landmarks in each sequence and the distance between each landmark is 0.75s.

The weight between sensory inputs and the place cell layers ($\mathbf{W}_{SP}$ and $\mathbf{W}_{PS}$) are learned by pseudoinverse learning rule [6] in an online manner [8] while exploring the environment as follows:

$$\mathbf{W}_{SP} = \mathbf{S} \cdot \mathbf{P}^{\dagger}, \tag{2}$$

$$\mathbf{W}_{PS} = \mathbf{P} \cdot \mathbf{S}^{\dagger}, \tag{3}$$

where $\mathbf{S} \in \mathbb{R}^{N_s \times N}$ and $\mathbf{P} \in \mathbb{R}^{Np \times N}$ denote sensory patterns and place patterns respectively, and $\dagger$ indicates the pseudoinverse.

Given the sensory input at time t, $\mathbf{s}_t$, its corresponding place cell and grid cell activations are computed as follows:

$$\mathbf{p}_t = \mathsf{sign}(\mathbf{W}_{PS} \cdot \mathbf{s}_t), \tag{4}$$

$$\mathbf{g}_t = \mathsf{CAN}(\mathbf{W}_{GP} \cdot \mathbf{p}_t). \tag{5}$$

where $\mathsf{CAN}(\cdot)$ represents the continuous attractor recurrence in the grid layer that is implemented using a module-wise winner-take-all dynamics (one-hot for each grid cell).

The grid cell activation is modulated by velocity signals, an action predicted by CTRNN in our case. This emulates the path integration, wherein the activation index of each grid cell module is shifted based on the direction of the action to deduce the subsequent grid state. Upon acquiring the subsequent grid code, $\mathbf{g}_{t+1}$, its related place code, $\mathbf{p}_{t+1}$, is correlated with the sensory input represented by $\mathbf{s}_{t+1}$.

B Details of Behavioral Study

All human experiments were approved by the Committee on the Use of Humans as Experimental Subjects at our organization. Subjects conducted six sessions of which each one takes 60 minutes on a different day. There are three 9-landmark sequences where the interval between each landmark is 0.75s as shown in Figure 5. Subjects press the left and right arrow keys on the keyboard to navigate toward the correct target landmark from the start landmark. If they stop at the target landmark within 0.25s error range, it is considered as correct. If a subject corrects at 15 pairs among the recent 20 pairs or tries more than 360 times in each case, the next case is started. Table 1 illustrates a sequence of case in each session. Each landmark sequence is trained as NTS for two consecutive sessions and keeps testing on the following sessions as MNAV. In the last session, subjects were tested on entire landmark sequences both NTS and MNAV settings.

C Experimental Details

We model a landmark as a randomly generated 384-dimensional vector with the same size of interval that is modeled as $\mathbf{0}$-vector. As an agent moves in one direction, the input vector is shifted to 64

Table 1: Sequence of task in each session for behavioral study. Each landmark sequence is trained for two consecutive sessions and keeps testing

Session	Landmark Sequence	Type	Pairs
session 1	Sequence 1	NTS	Training
	Sequence 1	MNAV	Training
session 2	Sequence 1	NTS	Training
	Sequence 1	MNAV	Training + New
	Sequence 1	NTS	Training + New
session 3	Sequence 2	NTS	Training
	Sequence 2	MNAV	Training
	Sequence 1	MNAV	Training + New
session 4	Sequence 2	NTS	Training
	Sequence 2	MNAV	Training + New
	Sequence 1	MNAV	Training + New
	Sequence 2	NTS	Training + New
session 5	Sequence 3	NTS	Training
	Sequence 3	MNAV	Training
	Sequence 1	MNAV	Training + New
	Sequence 2	MNAV	Training + New
session 6	Sequence 3	NTS	Training
	Sequence 3	MNAV	Training + New
	Sequence 1	MNAV	Training + New
	Sequence 2	MNAV	Training + New
	Sequence 3	NTS	Training + New
	Sequence 1	NTS	Training + New
	Sequence 2	NTS	Training + New

dimensions toward the direction, which means that the image-to-image distance is 12 steps. Regarding Hippocampal-MESH, the grid periods are 11, 12, 13 and the number of place cells is 400. The dimension of hidden states of CTRNN is 256 and its decaying factor is 0.9. The hidden states pass ReLU activation followed by one fully-connected layer to predict action (left, right, stop). We train the models for 2000 epochs for each environment and the maximum length of each episode is 100. We use Adam [2] with learning rate 0.001.