

Lecture 13: Neuroscience Theory for Navigation

Scribes: Shinwoo Choi, Yahav Avigal, Zhe Cao, Ilija Radosavovic

13.1 Navigational System In The Brain

This section summarizes the presentation on “*Microstructure of a spatial map in the entorhinal cortex*” [7] by Albert J. Qü.

13.1.1 Introduction

Brief Outline of hippocampus function in navigation. Fig. 13.1 [8] shows interesting relationships between recurrent nature for the Entorhinal Cortex in hippocampus (EC-HPC) system and the recurrent architecture of common navigational system. They both are very similar to each other. It is also interesting that the way theta band works is strongly reminiscent of CPU clocks

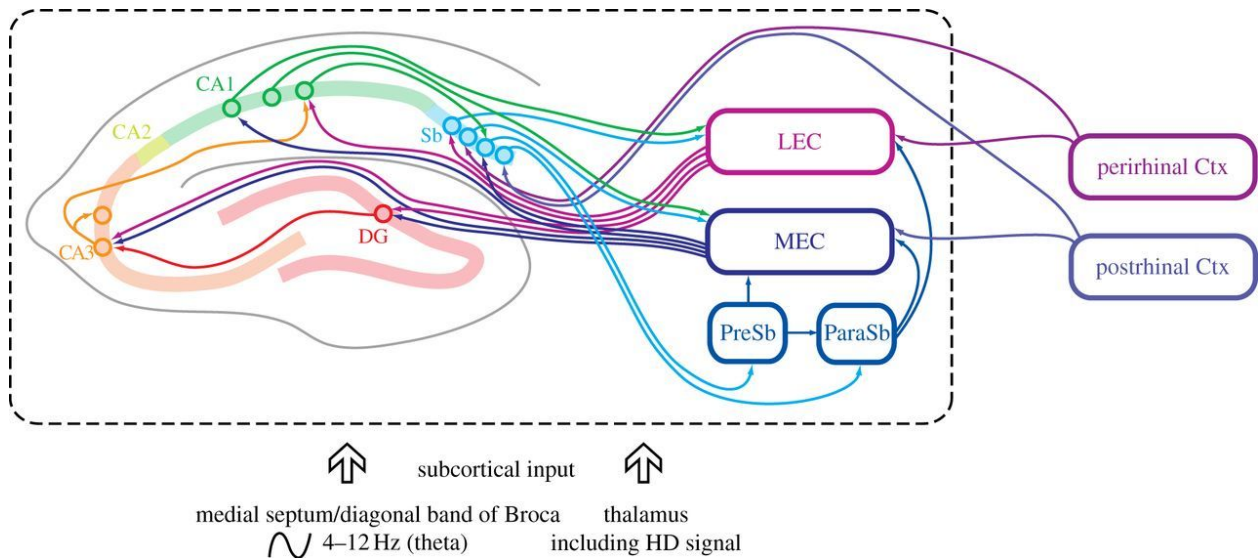


Figure 13.1: Schematic overview of major anatomical pathways in the hippocampal formation of the rat

Entorhinal Cortex in hippocampus system. EC-HPC system plays a plethora of roles in navigation. Fig. 13.2 shows that rats forage for food in a square-shaped environment. Whenever a spike occurs, the location of the animal is recorded. The heat map shows the spike rate of a given cell as a function of location. The color scale ranges from no activity in dark blue to the maximum rate in red [15] and Polar plots indicating strong directional tuning of firing rate in the head direction cells [16]

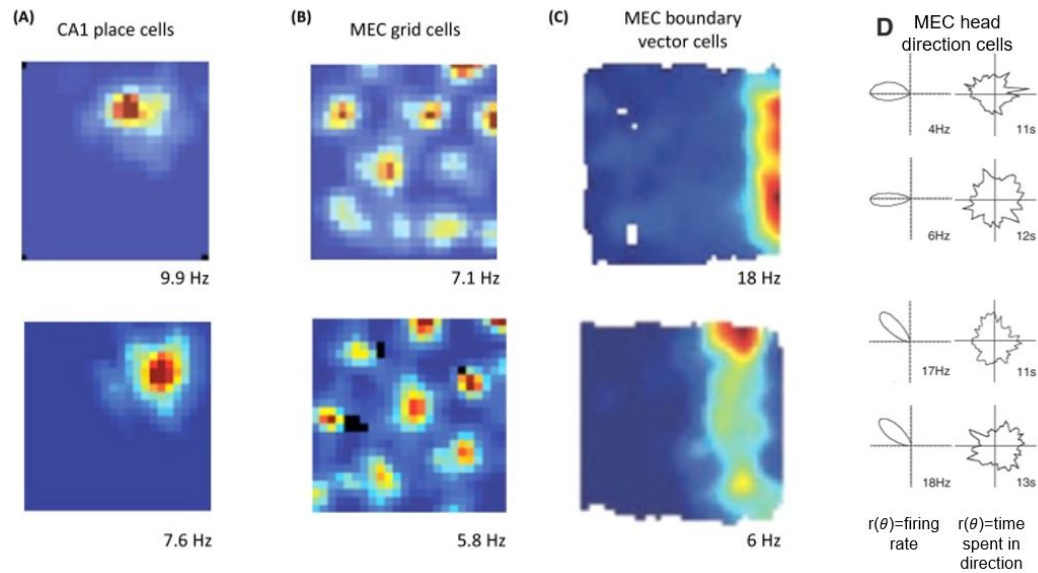


Figure 13.2: Firing Properties of Cell Types

13.1.2 Microstructure of a spatial map in the entorhinal cortex

Brief overview of the grid cell.

- Grid cells are spatially tuned in hexagonal patterns
- Grid cells are topologically arranged in the MEC
- Receptive fields (“navigational map”) for individual cells anchored to spatial/environment cues
- Robust properties of the navigational map
 - Non-reliant on visual cues/lighting conditions
 - Fast rate of convergence/stability

Navigational map of a single cells. Fig. 13.3 shows that brief overview of the firing fields of grid cells which have a repetitive triangular structure. Each cell forms hexagonal patterns. Fig. 13.3 (b) shows that the firing field formed a grid of regularly tessellating triangles spanning the whole recording surface. All nodes of activity were sharply delineated from the background, although the individual peak firing rates varied. The regular nature of the activity distribution was verified by spatial autocorrelation analyses, which for all cells showed a tessellating pattern similar to that of the original rate maps. Fig. 13.3 (c) shows that Box plot showing distribution of angles ϕ_1 , ϕ_2 and ϕ_3 between the central peak of the autocorrelogram and the vertices of a hexagon defined by the nearest six peaks. These angles between two vertices are roughly 60 degrees.

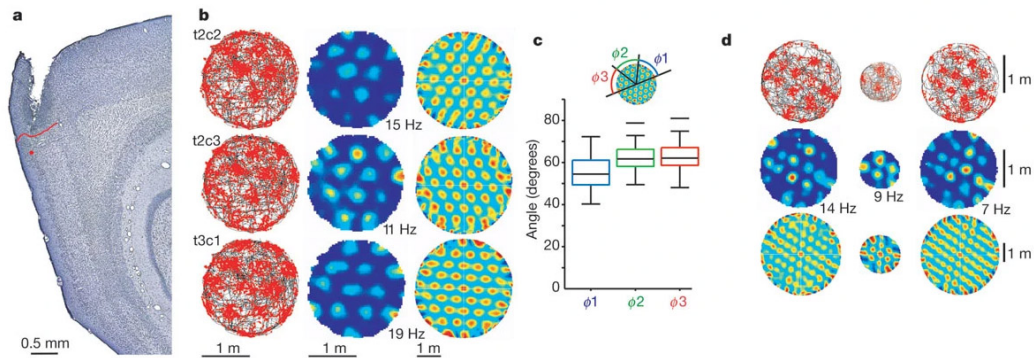


Figure 13.3: Firing fields of grid cells have a repetitive triangular structure.

Optimality of hexagonal grid representation. The hexagonal tiling is the densest way to arrange circles in two dimensions.

- Honeycomb conjecture: it states that a regular hexagonal grid or honeycomb is the best way to divide a surface into regions of equal area with the least total perimeter.

$$\limsup_{r \rightarrow \infty} \frac{\text{perimeter}(C \cap B(0,r))}{\text{area}(C \cap B(0,r))} \geq \sqrt[4]{12}$$

(C is the union of sets of bounded connected components with finite graph complement)

- Densest circle packing: In the two-dimensional Euclidean plane, Joseph Louis Lagrange proved in 1773 that the highest-density lattice packing of circles is the hexagonal packing arrangement, in which the centres of the circles are arranged in a hexagonal lattice.

$$\frac{3\pi}{4} D^2 / \frac{3\sqrt{3}}{2} D^2 = \frac{\pi\sqrt{3}}{6} \approx 0.9069 \quad (D: \text{diameter of circles and side length for hexagons})$$

Grid cells are topographically organized. Grid cells in the dMEC showed a striking topographic organization. Grids recorded at the same electrode location shared a number of metric properties, including spacing, orientation (direction) and field size.

- Coordinate system in neuroscience (Fig 13.4).
 - Vertical direction: dorsal-ventral axis for z , grid for single cell decrease in spacing and size in Z of cell location, while grids for cells located in the same planes are phase shifted from each other
 - Horizontal direction: Anterior-Posterior axis for y and lateral axis for x .




Figure 13.4: Coordinate system in neuroscience.

- Evidences in spacing (Fig 13.5). Grid cell representation is quite dense in dorsal, and as you go more and more to ventral, the more spaced out between grid cells. Grid cell is also larger in ventral.




Figure 13.5: Evidences in spacing on dorsal and ventral.

Topological map in brain. Topological map seems one of the general options in brain architecture.

- For example in vision, there are ocular dominance columns or orientation tuning columns that arranged in 2d patterns (Fig 13.6(a)).
- We also have a similar topological map in brain for motor and sensory area with different parts of the body. (Fig 13.6(b)).
- This is very similar to modular design in computation.




Figure 13.6: Topological map in Vision, Motor and Sensory areas in brain.

Anchoring to the environment

- Grid cells have very high correlation in rotational position in an environment. For example, when the grid cells are placed in a certain environment with a certain stimulus in a direction, one with a certain direction and the other with a rotated direction. Grid cells produces very high correlation in their patterns. (Fig. 13.7)

- One possible explanation (perhaps less interesting) is that the brain simply establishes some anchor point and creates the map with that coordinate. It would be interesting, however, if the brain detects rotation, but gradually use feedback control to rotate back.

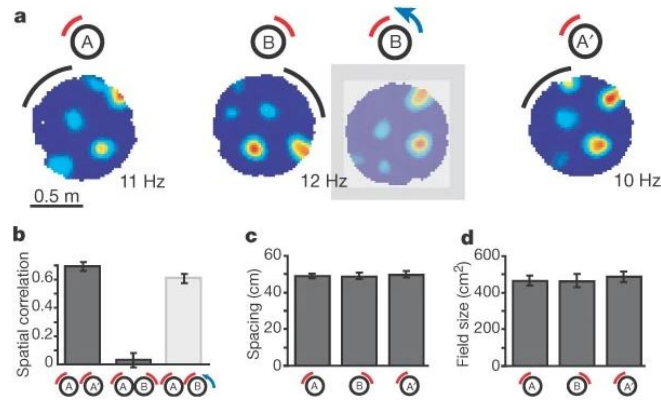


Figure 13.7: Grids are aligned to environment-specific landmarks.

Robust features(1) Invariance in light conditions

- Grid cell representation doesn't just disappear when you complete turn off the light.
- The grid cell architecture seems to be not merely a response to visual representation but functional without significant visual cues (Fig. 13.8)

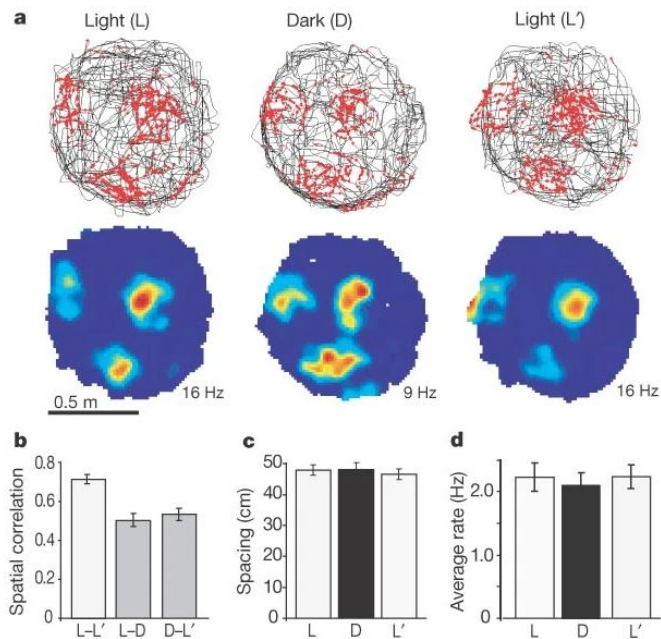


Figure 13.8: Grids persist in darkness.

Robust features(2) Rapid convergence and stability of map structure across time

- The grid cells are able to establish the stability of map structure across the time. When grid development in a novel environment, firing locations were mostly stable across the time rapidly (Fig. 13.9)

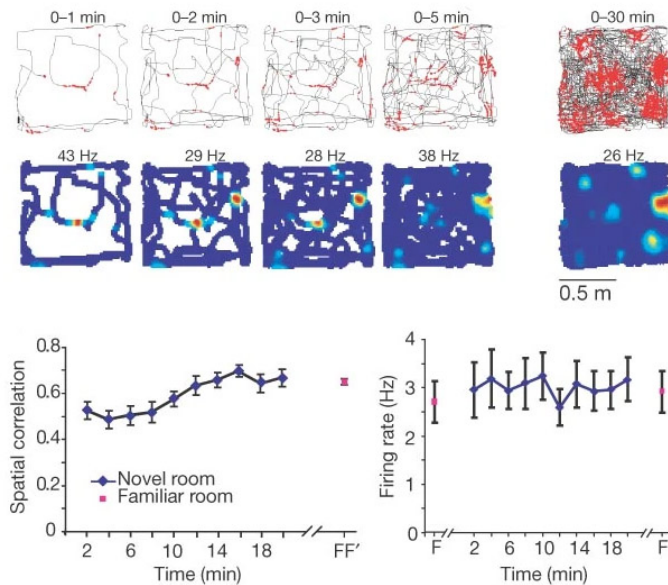


Figure 13.9: Grid structure of dMEC cells is expressed instantly in a novel environment.

Summary of grid cell

- Geometrically: optimal packing
- Anatomically: topological organization, phase shift and tuning of receptive field
- Functionally: stable and rapid formation in novel environment; not requiring visual context; anchoring to environmental cues

13.1.3 Other types of navigation related cells

Boundary detection cells. [17]

- Their representation is invariant to the room size or room shape. (Fig. 13.10 (a))
- Boundary cells react to environment remorph and such remorph is reversible (Fig. 13.10 (b))
- Mathematically these cells could signify:
 - some risk function giving punishment to reaching the end of the environment.
 - Or simply marking the end blocking the state transition probability.

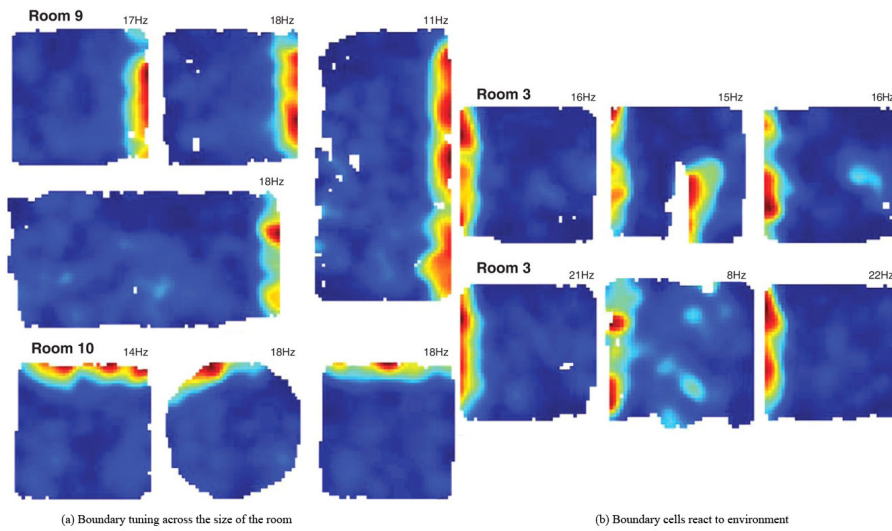


Figure 13.10: Boundary detection cells express proximity to boundaries in a number of environmental configurations.

Head-direction tuning Cells. [16]

- Selected subset of neurons not tuned for places nor displaying grid-like structure.
- The set of neurons instead tune to head orientation.
- Instead of placing a navigational map,

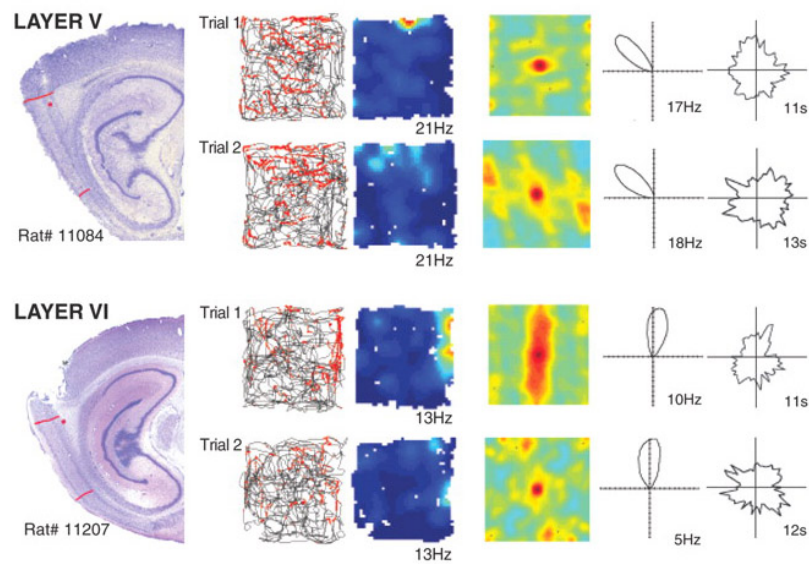


Figure 13.11: Head-direction cells in the MEC.

Place Cells. A type of a neuron within the hippocampus that fires when an animal enters a specific area in its environment. The place cells have interesting properties compared to the grid cells, such as local re-mapping, which happens when the animal enters a completely new area and the firing rate associated with a spatial coordinate changes, and global re-mapping where the place cells are associated with a new spatial coordinate after entering a new environment.

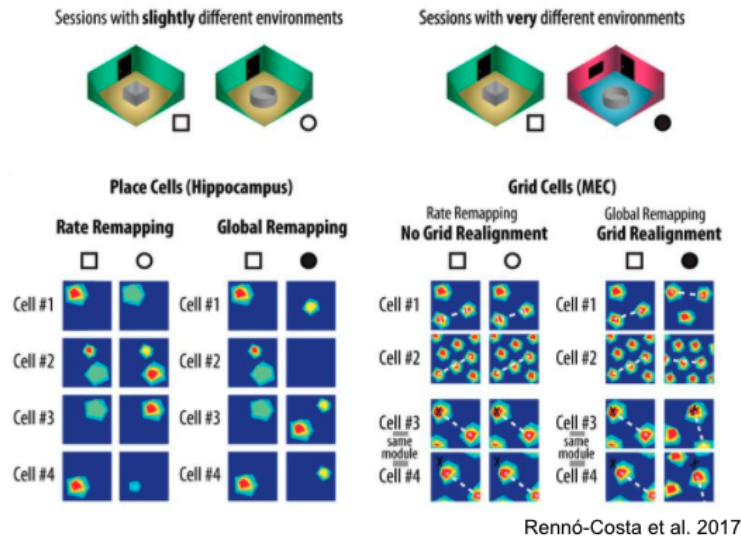


Figure 13.12: Place cells can re-map as a result of changes in the environment.

Extrom et al. [4] analyzed the behaviour of place cells in epileptic subjects by recording the signals of 317 neurons while the subjects navigated a virtual town, and learned that the cells fired when presented with specific locations even though the environment is virtual.



Figure 13.13: Extrom et al [4] experiments in a virtual environment

Dragoi et al. [3] investigated a phenomenon called "preplay" where temporal firing sequences reflecting the future order of place cell firing can happen during sleep or rest in some animals before they had any experience on long linear tracks (B in Fig.13.14), which indicates that temporal sequences of place cells are not necessarily caused by an ongoing external input-driven stimuli, but rather represent the internal organization of the hippocampal network.

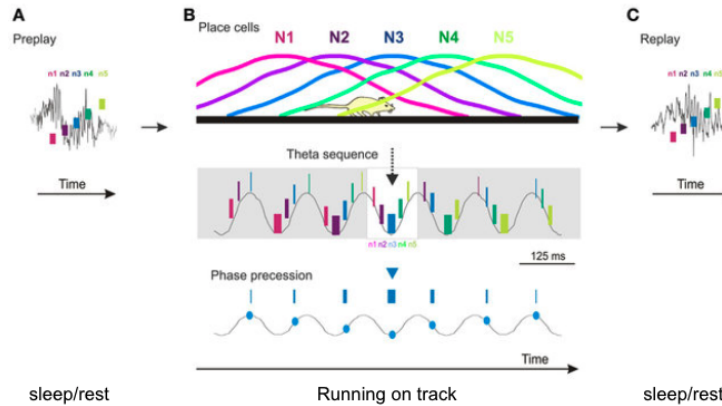


Figure 13.14: Dragoi et al [3] demonstrate "preplay".

Sanders et al. [15] discuss "theta sequences" where place cells fire during a theta cycle in order of the position of their place fields, as seen in Fig.13.15. In spite of the fact the subject animal moves very little during a theta cycle, each cell fires at the order of its position (where cell 1 represents position 1, cell 2 represents position 2, etc.) even if the corresponding location is ahead of the animal at the same moment in time, which can be thought of as anticipating or planning ahead. The firing patterns can be split to the first and second halves of the theta cycle, where in the first half they actively represent the current positions through positions coming from the path integration to the place cells and in the second half the grid cells integrate velocity vectors towards the "anticipated" locations (Fig.13.16).

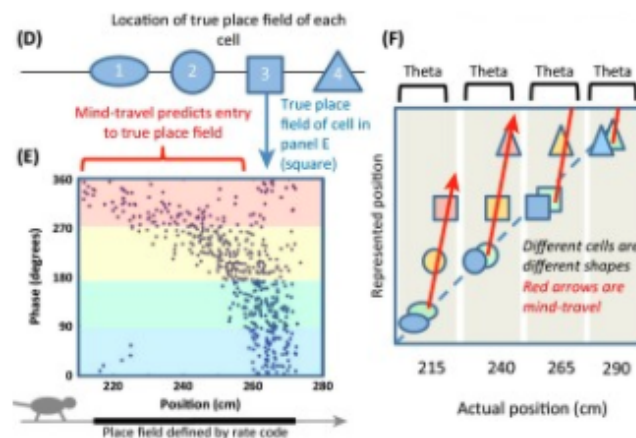


Figure 13.15: Sanders et al [15] show theta cycles.

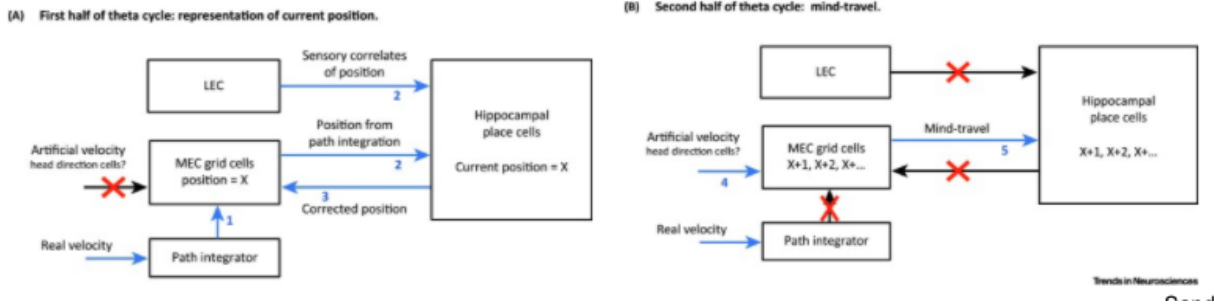


Figure 13.16: The firing patterns can be split to the first and second halves of the theta cycle.

13.2 Vector-based navigation using grid-like representations in artificial agents

This section summarizes the presentation on “*Vector-based navigation using grid-like representations in artificial agents*” [1] by Allen Shen.

The paper considers the problem of training deep reinforcement learning (RL) agents for navigation. While navigation was explored in prior work [14, 9, 12] it was still not at the level of other deep learning successes. The paper takes inspiration from mammals and aims to operationalize it within a computational framework.

13.2.1 Neuroscience theories

The proficiency of mammalian spatial behavior is underpinned by grid cells in the entorhinal cortex [7]. Grid cells are thought to provide multi-scale periodic representation that functions as a metric for coding space [6, 10]. This is believed to be critical for self-motion integration (path integration) [7, 6, 11] and planning direct trajectories to goals (vector-based navigation) [6, 5, 2].

13.2.2 Path integration

Path integration is basically tracking the position and direction of moving agent based on initial position using information about velocity. Medial entorhinal cortex (MEC) is the core of the path integration system. Specifically, place cells memorize past locations, head-direction cells sense movement and direction, and grid cells divide the spatial environment into a honeycomb hexagonal grid similar to the coordinate system on a map.

The paper trained an (Long Short-Term Memory network) LSTM using backpropagation through time for a path integration. The input to the network is the relative velocity, and the initial head direction. The output is the future location at each time step for both head-direction cells and Place cells. The LSTM which contains few linear layers is trained with cross-entropy loss and dropout.

As shown in Figure 13.17, remarkably, many units in the linear layer became like the grid cells found in the medial entorhinal cortex (MEC) of mammals – these cells had clearly defined firing fields that were arranged in a hexagonal pattern in both square and circular enclosures. Other cells had border- and head direction-related firing as well as conjunctive tuning, representing the full complement of functional cell-types found in the MEC.

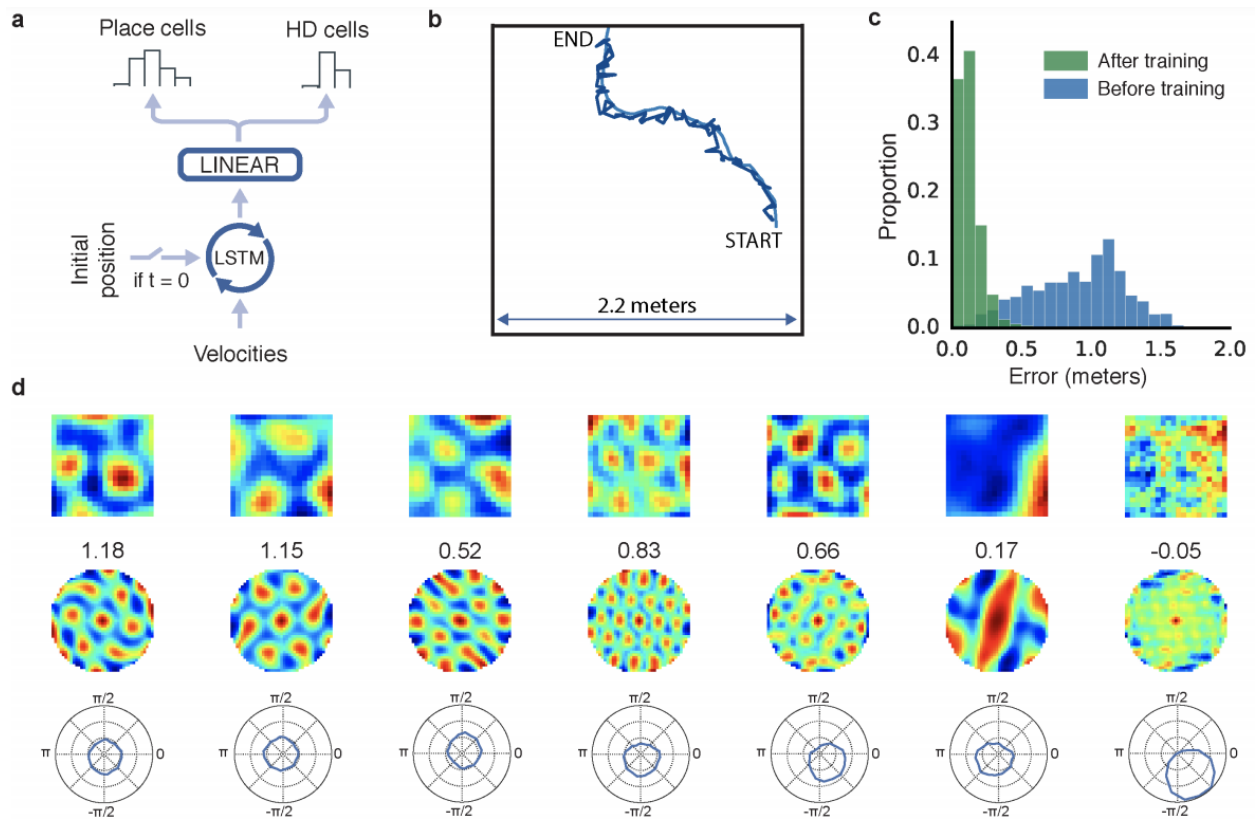


Figure 13.17: Entorhinal-like representations emerge in a network trained to path integrate. **a**, Schematic of network architecture (see Extended Data Figure 1 for details). **b**, Example trajectory (15s), self-location decoded from place cells resembles actual path (respectively, dark and light blue). **c**, Accuracy of decoded location before (blue) and after (green) training. **d**, Linear layer units exhibit spatially tuned responses resembling grid, border, and head direction cells. Ratemap shows activity over location (top), spatial autocorrelogram of the ratemap with gridness indicated (middle), polar plot show activity vs. head direction (bottom).

Importantly, the linear layer was subject to regularization, in particular dropout, such that 50% of the units were randomly silenced at each time step. Grid-like structure did not appear when dropout was not applied.

13.2.3 Vector-based navigation

Vector-based navigation enables the brain to calculate the distance and direction to a desired destination, “as the crow flies,” allowing animals to make direct journeys between different places even if that exact route had not been followed before. Entorhinal grid cells proposed to provide a Euclidean spatial metric, allowing for the calculation of goal-directed vectors and enabling animals to follow direct routes to a remembered goal. However, experimental evidence for the direct involvement of grid representations in goal-directed navigation is still lacking in existing work.

13.2.4 Navigation experiments

The authors use network with grid representations and train an agent for navigation tasks using deep RL. The Figure 13.18 shows the architecture of the grid cell agent which contains the grid network (light blue dashed) and the visual module (green dashed) and an actor critic learner (based on A3C [13]; dark blue dashed). The vision module, i.e., is a convolutional neural network that produces place and head direction cell activity patterns. The grid cell network of the agent was implemented as in the supervised learning set up except that the LSTM (“GRID LSTM”) was not initialised based upon ground truth place cell activations but rather set to zero. Velocity signals are perturbed with random noise as well as visual input to grid network. Interesting, input to grid network is only provided 5% of the time to mimic the imperfect visual observations. The output of the linear layer of the grid network, corresponding to the agent’s current location, was provided as input to the “policy LSTM”, a second recurrent network controlling both the agent’s actions and outputting a value function. Additionally, whenever the agent reached the goal, the “goal grid code” — activity in the linear layer — was subsequently provided to the policy LSTM during navigation as an additional input.

Conventional simultaneous localization and mapping (SLAM) techniques typically require an accurate and complete map to be built, with the nature and position of the goal externally defined. By contrast, the deep reinforcement learning approach described in this work has the ability to learn complex control policies end-to-end from a sparse reward, taking direct routes involving shortcuts to goals in an automatic fashion—abilities that exceed previous deep reinforcement learning approaches, and that would have to be hand-coded in any SLAM system.

13.2.5 Support for neuroscience theories

Overall, the results provide support for neuroscience theories. First, artificial grid cells resemble biological grid cells: (1) Percentage of grid cells (25%), head direction cells (10%), and border cells (8.7%) is similar to distribution in medial entorhinal cortex. (2) Directionality of the head direction units showed a six-fold symmetry, echoing the six-fold symmetry found in the human brain during active navigation.

There are also evidence that grid cells enable vector-based navigation: (1) Grid cell agent circled a fake goal grid code when real goal was not present (similar to rodent behavior). (2) Withholding goal grid code from policy LSTM led to poor performance. (3) Decoded Euclidean distance and allocentric goal direction represented more strongly in grid cell agent compared to place cell agent. (4) Silencing grid-cells resulted in worse performance and impaired navigation compared to silencing non-grid cells. (5) Place cells provide a robust representation of self-location but are not thought to provide a substrate for long range vector calculations.

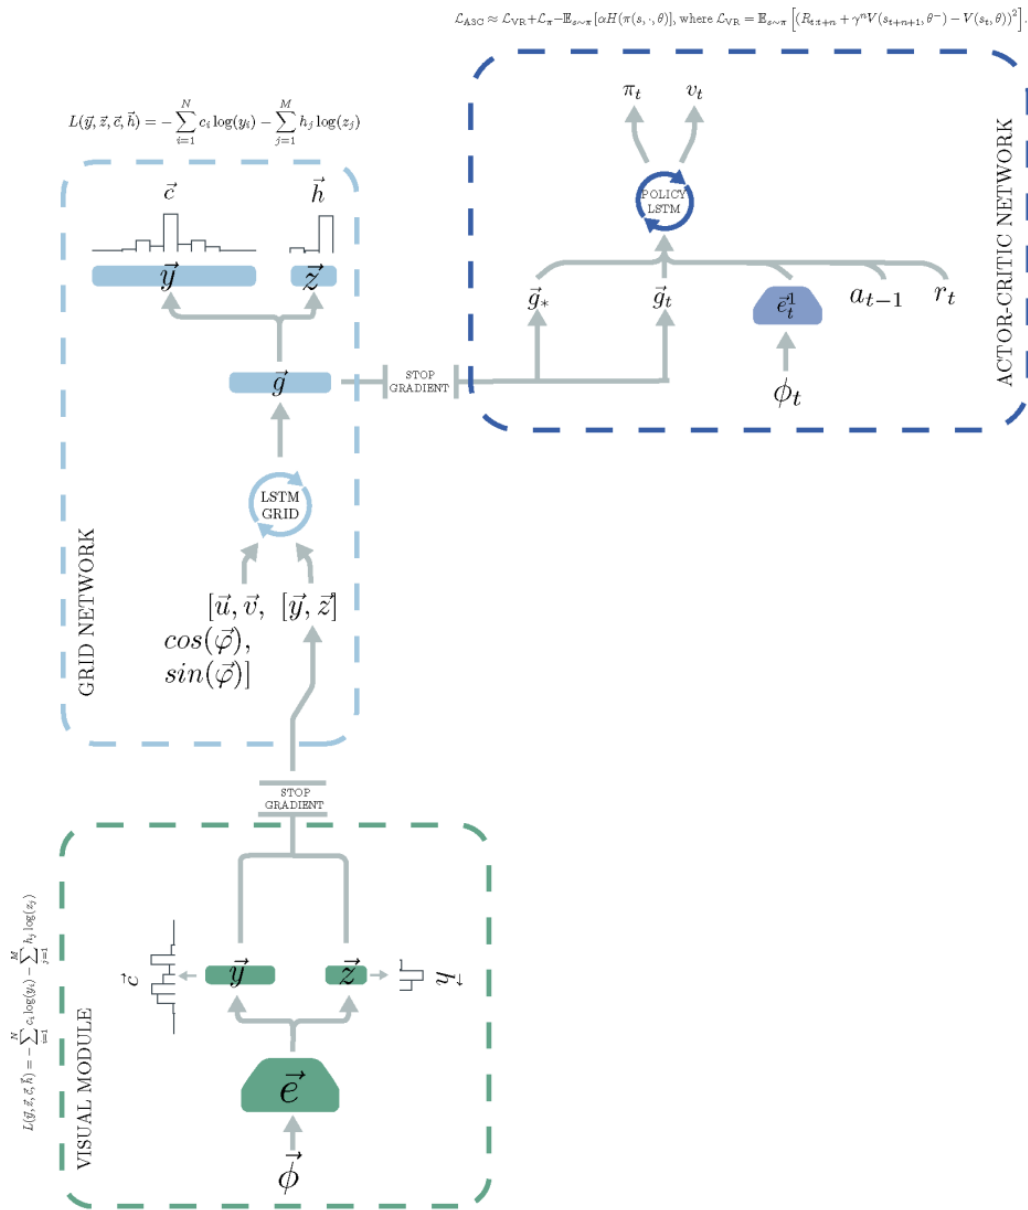


Figure 13.18: The architecture of the supervised network (grid network, light blue dashed) was incorporated into a larger deep RL network, including a visual module (green dashed) and an actor critic learner (based on A3C [13]; dark blue dashed). In this case the supervised learner does not receive the ground truth c_0 and h_0 to signal its initial position, but uses input from the visual module to self-localize after placement at a random position within the environment. Visual module: since experimental evidence suggests that place cell input to grid cells functions to correct for drift and anchor grids to environmental cues, visual input was processed by a convolutional network to produce place cell (and head direction cell) activity patterns which were used as input to the grid network. The output of the vision module was only provided 5% of the time to the grid network;

13.3 Discussions

Q1 (Professor Tomlin) When navigation is defined, is it just you know your destination and you have a map or a sort of learning the map as well?

A1 (Alber Qu) Good question! Throughout the presentation, I will hopefully show some of the evidences that it actually does a little of both. You can see a localization of the current location. and also it can be seen that there might be a searching or planning for the future motions.

Q2 (Professor Ma) For different directions that they encode the information with different fequency of firing?

A2 (Alber Qu) Yes, for example, if I am going to the north, then there is a north touring cell that is firing very hard.

Q3 (Professor Ma) So, you're saying that there is a sort of carrying signal? I mean that there are some kinds of periodic waves that does those firing pattern were carried down different phases? and those waves are physical?

A3 (Alber Qu) Yes. you can measure them. One of the results is that when people place these electrodes in the brain, you will be able to measure that. Its functioning as a clock across the different regions in the brain. In my understanding, the field is still developing people all have completing hypothesis.

Q4 (Kshitij Kulkarni) Do we know the evolutionary history of these parts of the brain / why they might be particularly attuned to recognizing high-frequency regions vs. body parts, etc?

A4 (Alber Qu) To my knowledge to answer the evolutionary history is very difficult because we don't have access to that. However there are some studies different levels like birds, as you go higher in a evolutionary tree, it might be able to develop some correspondences but as you go lower down it becomes harder and harder to record the brain activity.

Q5 (Professor Ma) When animal goes from familiar place to novel environment, what do they adjust or do with place cells?

A5 (Alber Qu) There are two possible things that might happen for the Place cells. In general, when animal is placed in the new environment, some cells that previously encoded with some certain x and y, they still fire but relatively insignificant(decreased firing rate or some differences in firing rate) in the same x and y in another environment. However, there is a also global remapping in the completely new environment. Firing takes place in completely different x and y (actually x' and y')

References

- [1] A. Banino, C. Barry, B. Uribe, C. Blundell, T. Lillicrap, P. Mirowski, A. Pritzel, M. J. Chadwick, T. Degris, J. Modayil, et al. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 2018.
- [2] D. Bush, C. Barry, D. Manson, and N. Burgess. Using grid cells for navigation. *Neuron*, 2015.
- [3] G. Dragoi. Internal operations in the hippocampus: single cell and ensemble temporal coding. *Frontiers in systems neuroscience*, 2013.
- [4] A. D. Ekstrom, M. J. Kahana, J. B. Caplan, T. A. Fields, E. A. Isham, E. L. Newman, and I. Fried. Cellular networks underlying human spatial navigation. *Nature*, 2003.
- [5] U. M. Erdem and M. Hasselmo. A goal-directed spatial navigation model using forward trajectory planning based on grid cells. *European Journal of Neuroscience*, 2012.
- [6] I. R. Fiete, Y. Burak, and T. Brookings. What grid cells convey about rat location. *Journal of Neuroscience*, 2008.
- [7] T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, and E. I. Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 2005.
- [8] T. Hartley, C. Lever, N. Burgess, and J. O'Keefe. Space in the brain: how the hippocampal formation supports spatial cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 2014.
- [9] T. D. Kulkarni, A. Saeedi, S. Gautam, and S. J. Gershman. Deep successor reinforcement learning. *arXiv:1606.02396*, 2016.
- [10] A. Mathis, A. V. Herz, and M. Stemmler. Optimal population codes for space: grid cells outperform place cells. *Neural computation*, 2012.
- [11] B. L. McNaughton, F. P. Battaglia, O. Jensen, E. I. Moser, and M.-B. Moser. Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience*, 2006.
- [12] P. Mirowski, R. Pascanu, F. Viola, H. Soyer, A. J. Ballard, A. Banino, M. Denil, R. Goroshin, L. Sifre, K. Kavukcuoglu, et al. Learning to navigate in complex environments. *arXiv:1611.03673*, 2016.
- [13] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, 2016.
- [14] J. Oh, V. Chockalingam, H. Lee, et al. Control of memory, active perception, and action in minecraft. In *International Conference on Machine Learning*. PMLR, 2016.
- [15] H. Sanders, C. Rennó-Costa, M. Idiart, and J. Lisman. Grid cells and place cells: an integrated view of their navigational and memory function. *Trends in neurosciences*, 2015.
- [16] F. Sargolini, M. Fyhn, T. Hafting, B. L. McNaughton, M. P. Witter, M.-B. Moser, and E. I. Moser. Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, 2006.
- [17] T. Solstad, C. N. Boccara, E. Kropff, M.-B. Moser, and E. I. Moser. Representation of geometric borders in the entorhinal cortex. *Science*, 2008.