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

Andrea Banino, Caswell Barry, Benigno Uria, Charles Blundell, Timothy Lillicrap, Piotr Mirowski, Alexander Pritzel, Martin Chadwick, Thomas Degris, Joseph Modayil, Greg Wayne, Hubert Soyer, Fabio Viola, Brian Zhang, Ross Goroshin, Neil Rabinowitz, Razvan Pascanu, Charlie Beattie, Stig Petersen, Amir Sadik, Stephen Gaffney, Helen King, Koray Kavukcuoglu, Demis Hassabis, Raia Hadsell, and Dharshan Kumaran

Presented by Roman Ring

Grid Cells

- Neurons that fire up in a specific pattern
  - *Hexagonal activations when exploring environment*
- Located in the entorhinal cortex
  - *“Interface” for neo-cortex and hippocampus*
  - *Used for memory, navigation, time perception*
- Crucial discovery in multiple fields
  - *2014 Nobel Prize in Physiology or Medicine*

Grid Cells in Rodents

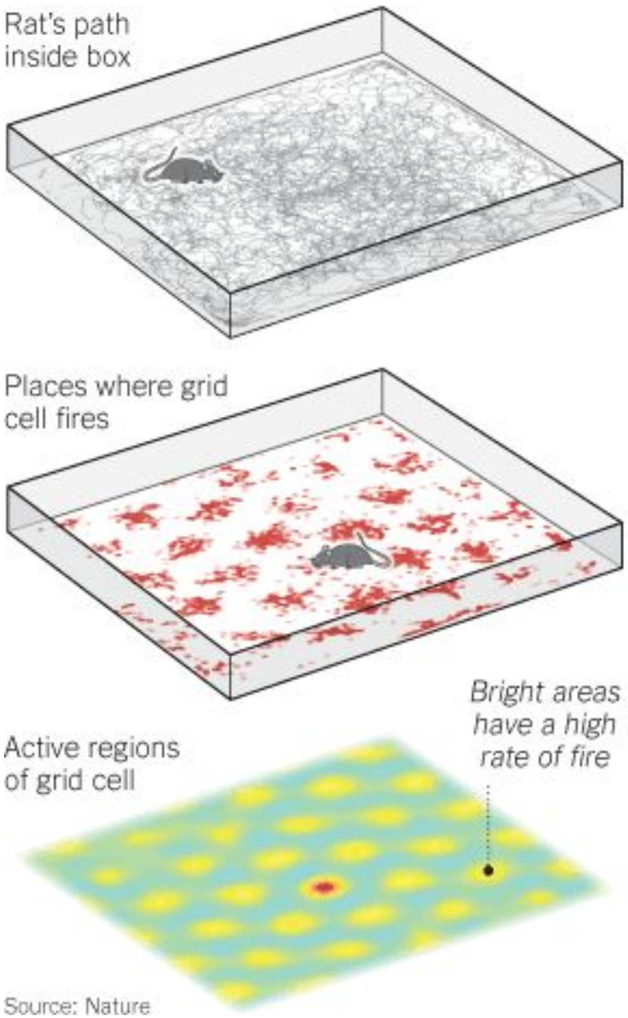

Source: Nature

Mental Maps

Researchers are studying how brain cells in the entorhinal cortex help rats and other mammals build maps of the environment.

A RANDOM WALK
At left, gray lines show the rat's path as it moves around a box eating pieces of food.

IMPOSING A PATTERN
Grid cells in the rat's entorhinal cortex fire when the rat moves through certain locations. The firing pattern of a single grid cell is marked here with dots. Groups of dots form a hexagonal grid, and the firing pattern persists even in darkness, when the rat cannot see where it is.

GRID CELLS
The grid cells seem to form a map of the local environment. Each grid cell, like the one enlarged at left, fires in a hexagonal pattern that helps the rat track where it is in space. Grid cells are thought to be involved with pathfinding, dead reckoning and the formation of mental maps.

THE NEW YORK TIMES

Source: New York Times

Vector-based Navigation

- Determine position and direction
  - *Including very noisy observations (e.g. closing eyes)*
- Determine distances between objects
  - *Utilizing Euclidean-like metrics*
- Efficiently navigate in an environment
  - *Find approximate shortest paths "in real time"*

Artificial Agents

- Brain is the best "computer" we know
  - *Very complex & massively parallel processing unit*
- Difficult to investigate hypotheses
  - *Even with mice experiments are very expensive*
- AI agents as an approximation
  - *Relatively fast proof of concept tests (?)*

Intuitive Interpretations

- Grid cells are used as a map of the world
  - *Special neurons fire up in a hexagonal grid*
  - *Position self in the environment*
- Extending to vectors for navigation
  - *Find shortest paths with Euclidean distance*
- Artificial agents used as a brain model
  - *Possible alternative to expensive experiments*

Research Objectives

- Show similarity between ANN and NN
  - *In the context of grid cells: positional reasoning*
- Train an RL agent that utilizes grid cells
  - *Show that it performs better with grid cells*
- Extrapolate results to neuroscience
  - *Show grid cells used for vector-based navigation*

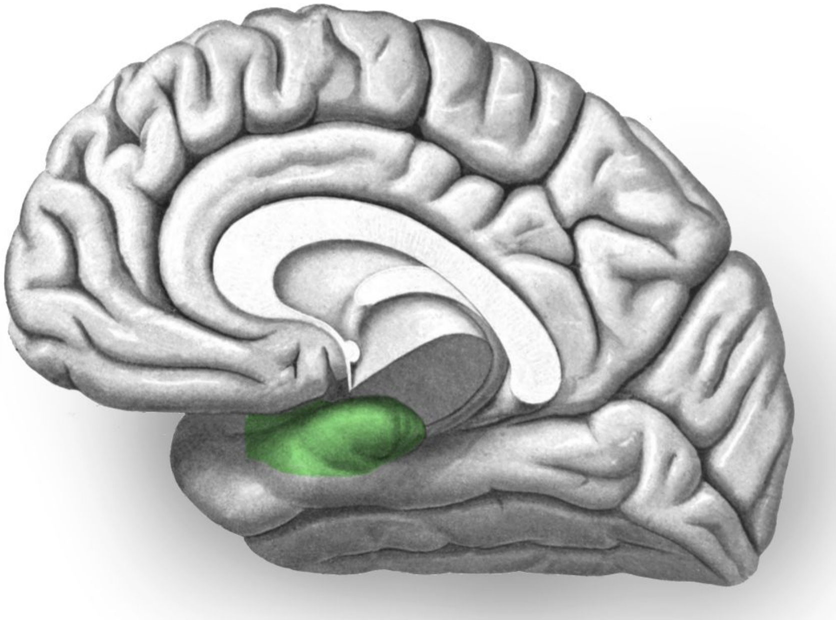

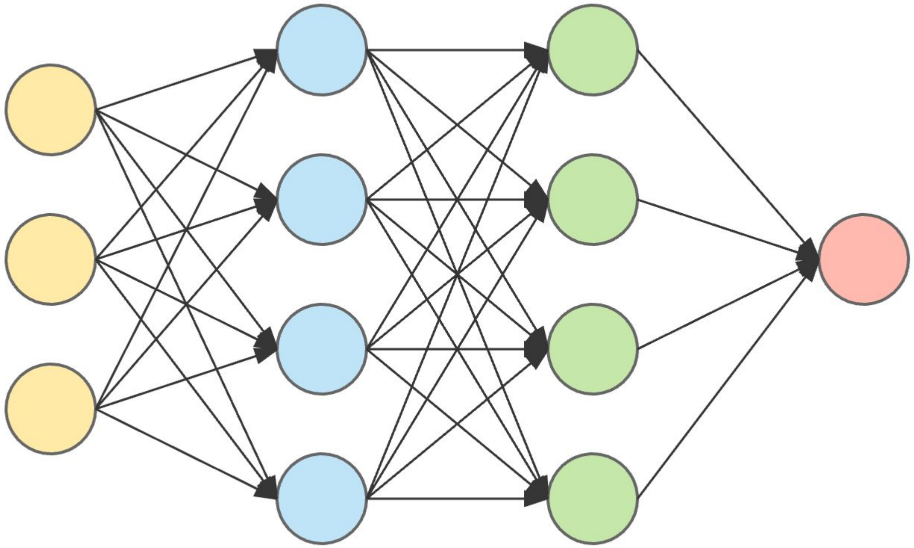

Showing NN & ANN similarity

Based on activations in rodent brains

Source: Radiopaedia, Towards Data Science

Artificial Neural Networks

- Non-linear function approximators
    - *Approximate any function with some assumptions*
- Very crude model for the actual neurons
    - *Still using McCulloch–Pitts model from the 50s*
- Reccurent ANNs for temporal data
    - *Hidden state that can add info or "forget" it*

Objective: Show Similarity

- Generate training data from rat models
  - *Input: place cells + head direction*
  - *Output: next place cell + head direction*
- Train an ANN on the data (supervised)
  - *Recurrent + Linear + Softmax*
- Compare ANN outputs to rat models
  - *Both visually (via spatial maps) and statistically*

ANN Architecture

- Small Recurrent ANN
  - *LSTM with 128 unit hidden cell*
  - *Linear layer with 512 units (no activation!)*
  - *Output heads: offset vector and angle activations*
  - *Probability distribution via softmax*
- No special tricks involved
  - *Important since goal is to show similarity*

ANN Architecture

$$\mathcal{L}(\vec{y}, \vec{z}, \vec{c}, \vec{h}) = -\sum_{i=1}^{N} c_i \log(y_i) - \sum_{j=1}^{M} h_j \log(z_j)$$

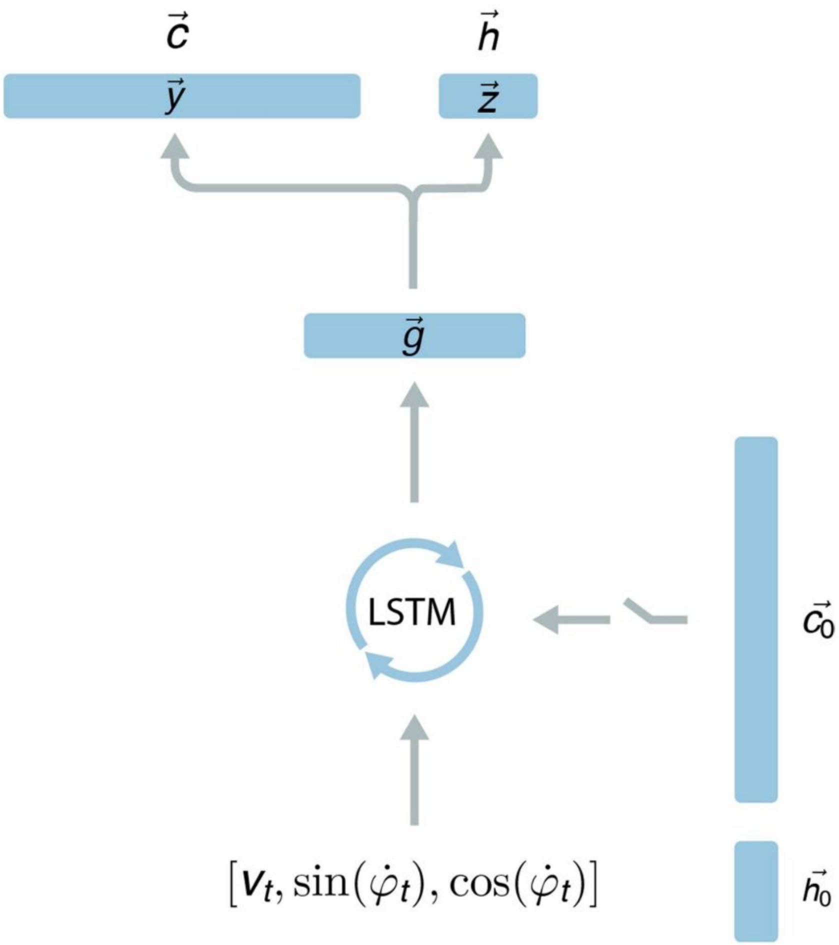

Generating Training Data

- Position centers uniformly sampled
  - *From some fixed box dimensions*
- Place activations via 2D Normal distribution
  - *With fixed mean and standard deviation*
- Direction angle via Von Mises distribution
  - *With fixed concentration parameters*

Emergence of Grid Cells

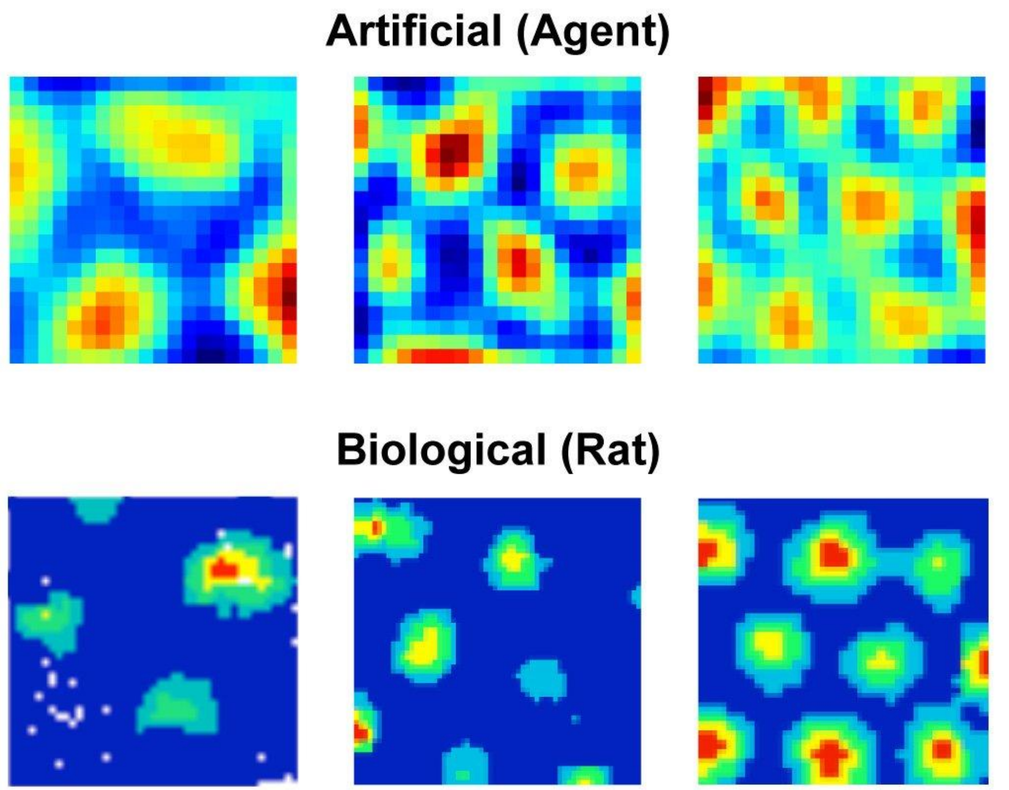

Our experiments with artificial agents yielded grid-like representations ("grid units") that were strikingly similar to biological grid cells in foraging mammals.

Source: DeepMind

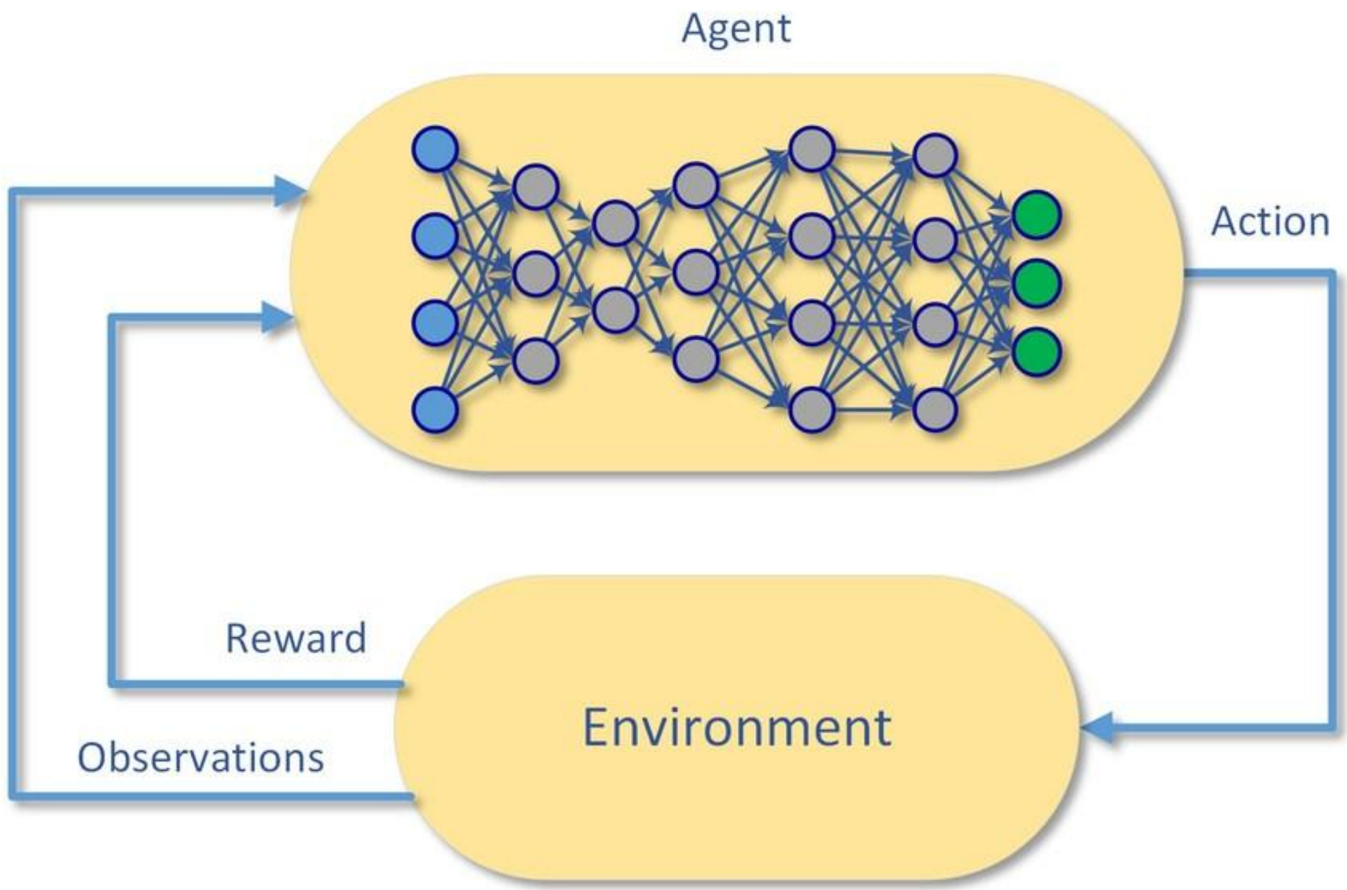

Training a DRL agent

A3C agent with grid cells navigating in a maze

Source: IEEE Communications Magazine

Deep Reinforcement Learning

- General framework for decision making
  - *Evaluate behavior with some reward signal*
- Policy is the expected output of training
  - *A mapping from states to actions*
- Algorithms differ by optimization target
  - *Value based, policy based, actor-critic hybrids*
  - *A3C: Asynchronous Actor-Critic Agent*

Objective: Train a DRL Agent

- Re-use network trained in previous step
  - *Add some noise in the inputs for generalization*
- Use “grid cells” network as an encoder
  - *Pipe into a much larger network*
- Evaluate ability to navigate in mazes
  - *Trained with A3C algorithm*

DRL Agent Architecture

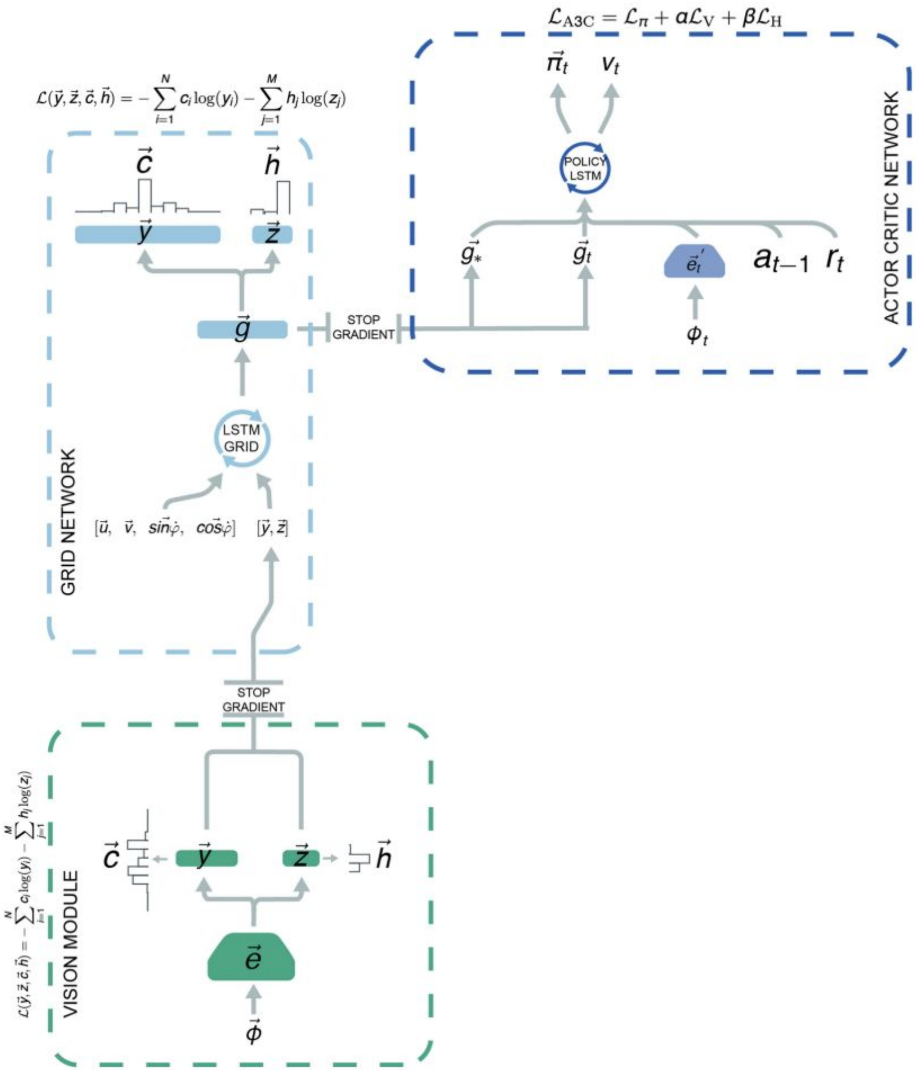

DRL Agent Navigating

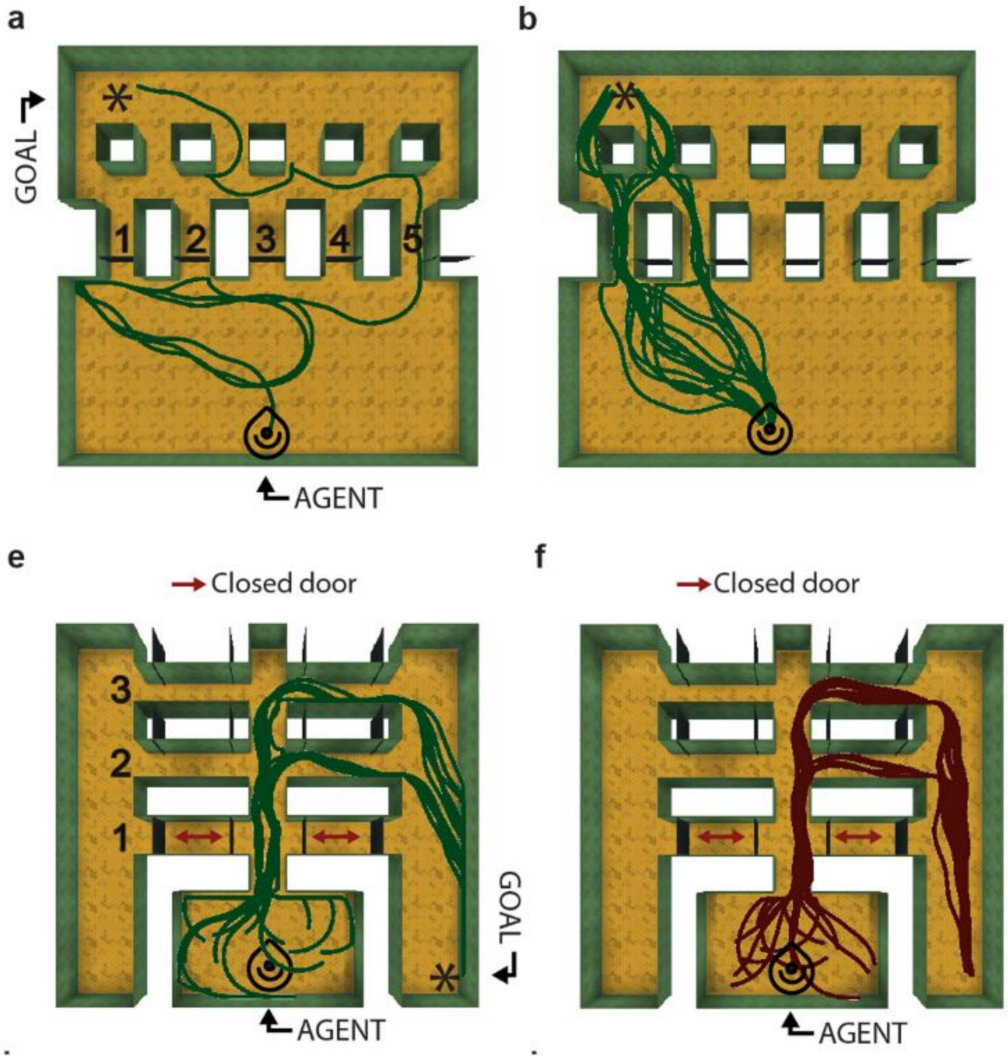

Effect of Grid Cells

Effect of Grid Cells

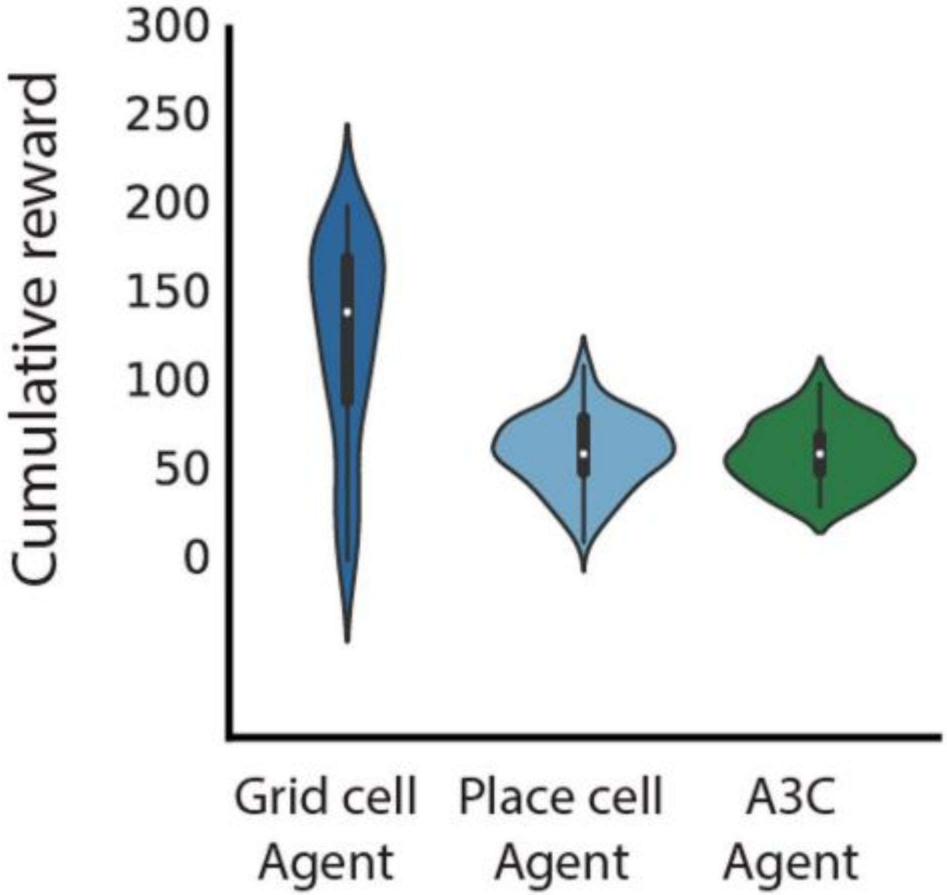

Extrapolating to Neuroscience

Emergence of vector-based navigation from grid cells

Objective: Extrapolate Results

- Grid cells emerged in artificial agents
- Disabling grid cells hindered navigation
  - *Agent performed significantly worse in known tasks*
  - *Agent failed to generalize to new tasks*
- We can extrapolate that grid cells used for vector-based navigation (in mammals)

Conclusion

- Shown emergence of grid cells in ANNs
  - *Similar activation patterns to rodent models*
- Utilized grid cells to train an RL agent
  - *Outperforming agents without grid cells*
- Shown use of vector-based navigation
  - *Assuming model is correct, can extrapolate that similar patterns emerge in actual grid cells*

Recap by DeepMind

References & Related Work

- Microstructure of a spatial map in the entorhinal cortex
- Crystals of the brain
- A Model of the Ventral Visual System Based on Temporal Stability and Local Memory
- Slowness and Sparseness Lead to Place, Head-Direction, and Spatial-View Cells
- Navigating with grid-like representations in artificial agents
- Emergence of grid-like representations by training recurrent neural networks to perform spatial localization