Modeling the function of episodic memory

Xiangshuai Zeng, Laurenz Wiskott, Sen Cheng

Institute for Neural Computation, Ruhr University Bochum

Introduction

It is critical to understand the function of episodic memory and how it is used in learning and decision making in order to really understand what episodic memory is. We postulate that episodic memory simply stores our past experiences, including important biographical events, such as our wedding day or an unforgettable birthday party, but also daily, mundane things, such as what we ate for breakfast or where we went the day before. Therefore, an episodic memory likely includes what, where, and when information, but it is not defined by this content.

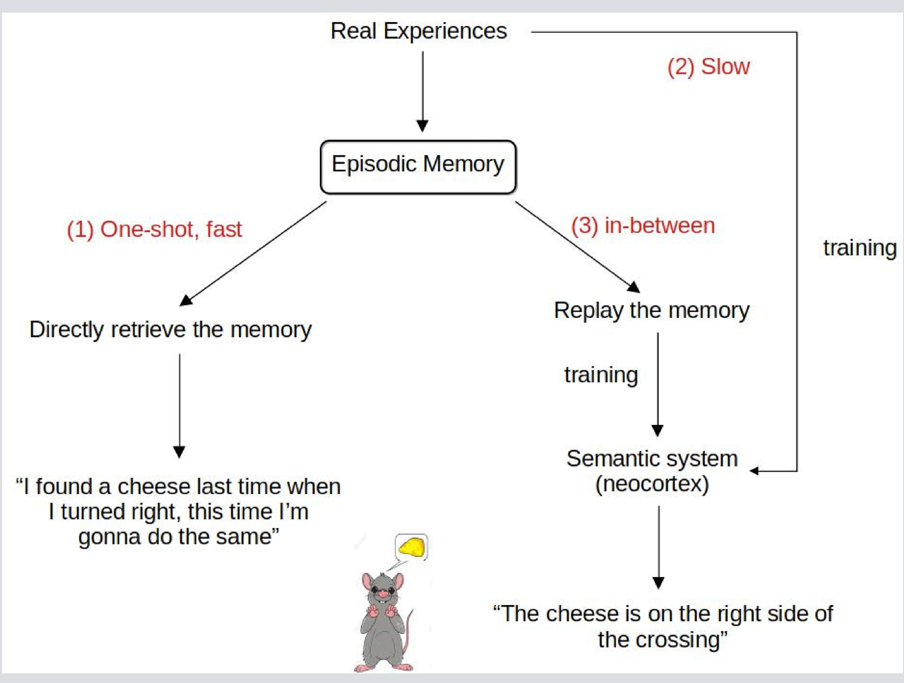

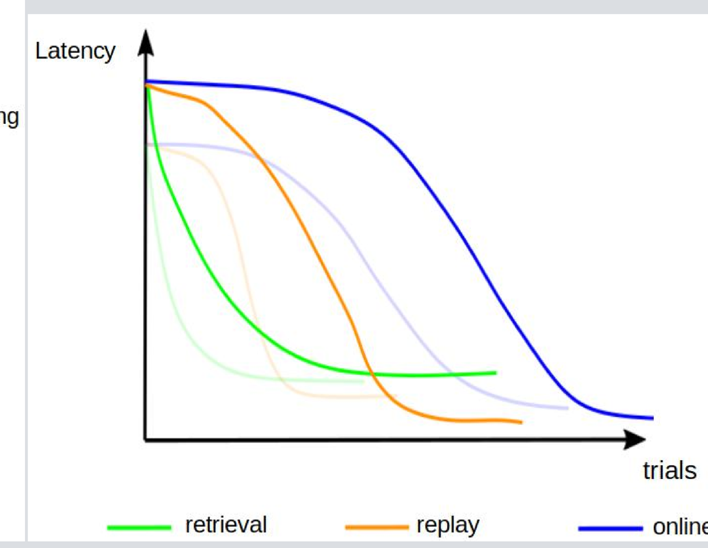

During learning and decision making, an agent could (1) retrieve an episodic memory and use the information contained in the memory directly for making a decision. (2) Alternatively, the animal could extract statistical regularities about the world, i.e., semantic information, while it is exploring the world. This, however, is slow. (3) To speed up learning, episodic memories could be replayed to train the semantic system.

The three learning paradigms have different learning curves. However, given enough trials, all three eventually learn the task, even without episodic memory. This might be one possible reason why some tasks are found to be hippocampally dependent in some experiments, but not in others.

Methods

To model the three learning paradiagms, we use reinforcement learning.

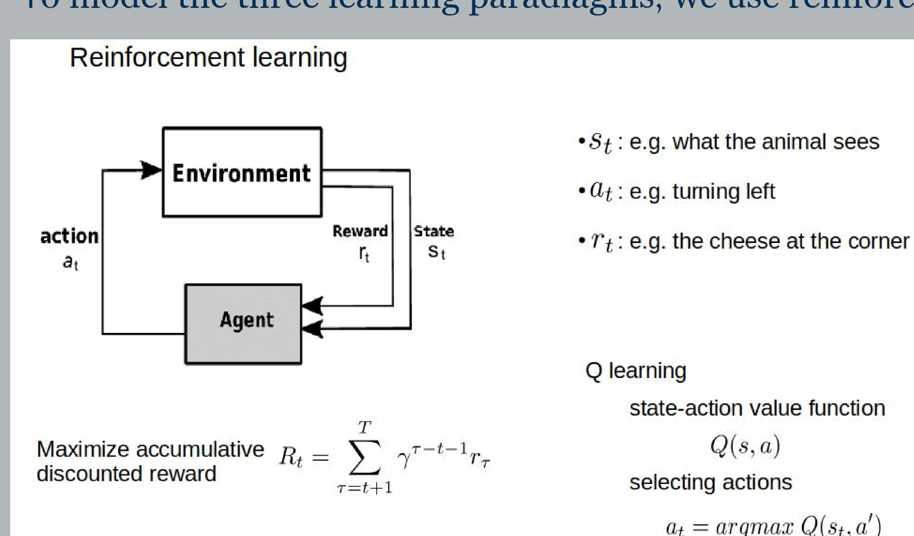

We employed three reinforcement learning algorithms: (1) model-free episodic control (memory retrieval) [1], (2) deep Q learning without experience replay (online learning), and (3) deep Q learning with experience replay (DQN) [2], and (3). We train the three agents in several environments with different task complexity and compare their learning curves.

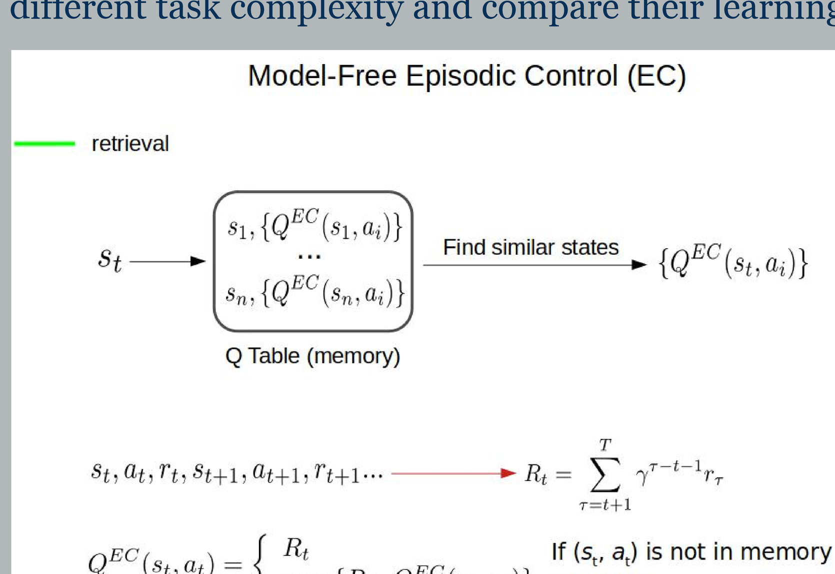

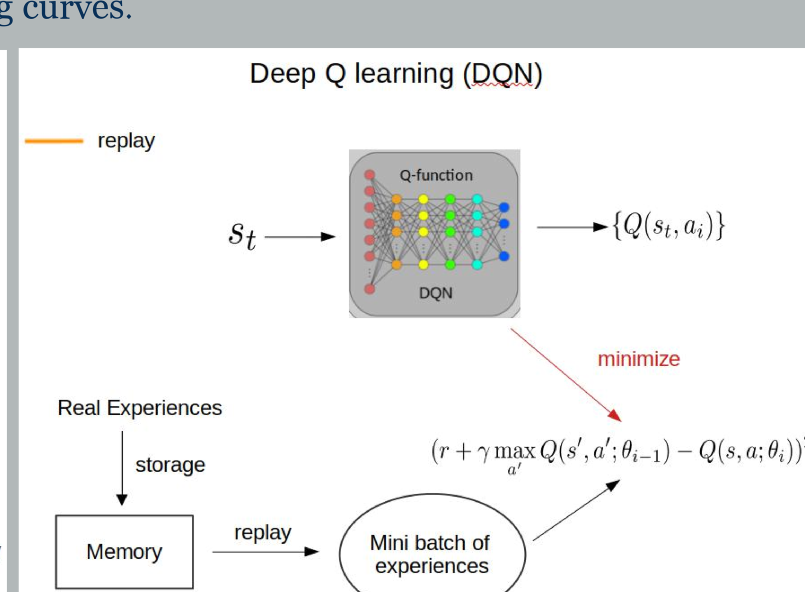

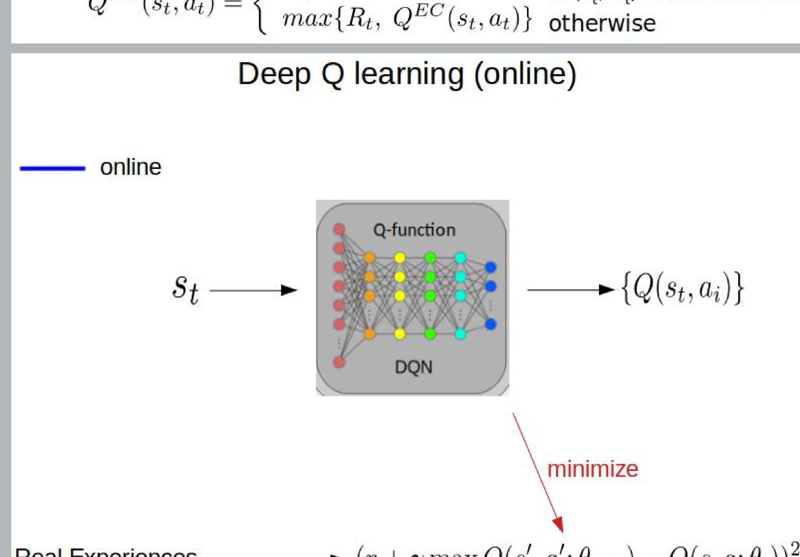

Results

We designed several mazes with the Unity simulator.

This is what the agent sees.

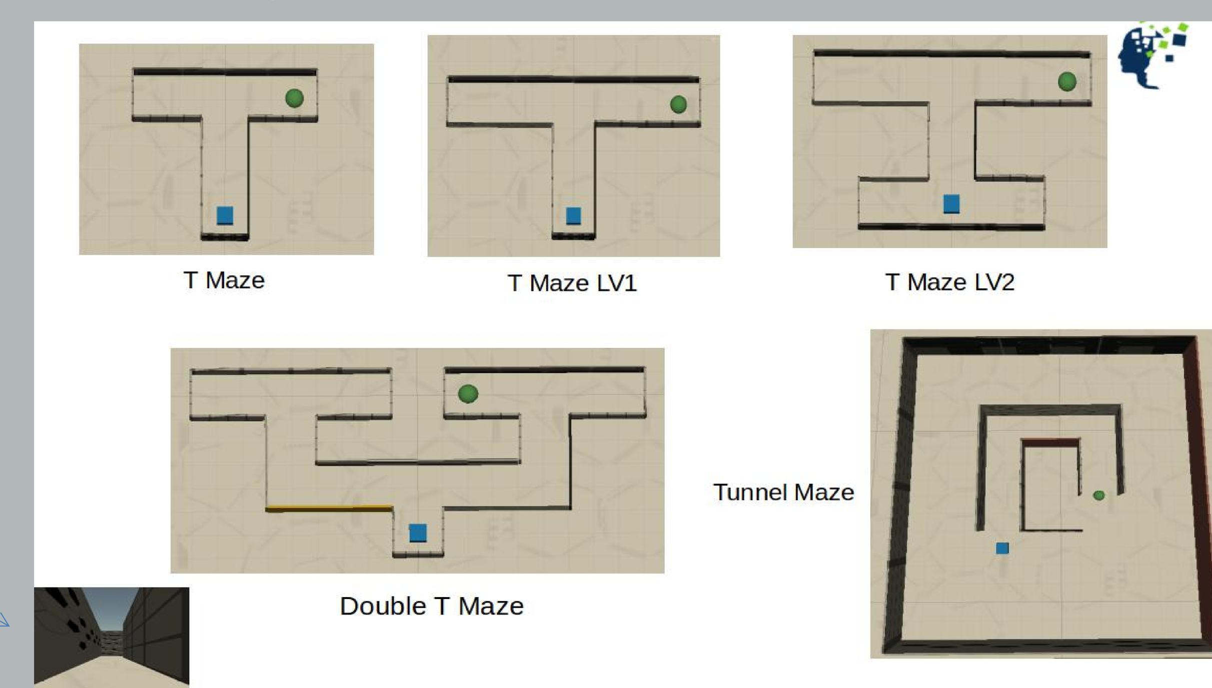

The blue square indicates an agent that can move up, down, left, right and rotate 90° clockwise and counter-clockwise. The green sphere indicates the location of a large reward. From left to right, top to bottom, the maze is becoming increasingly more complex.

Learning Curves

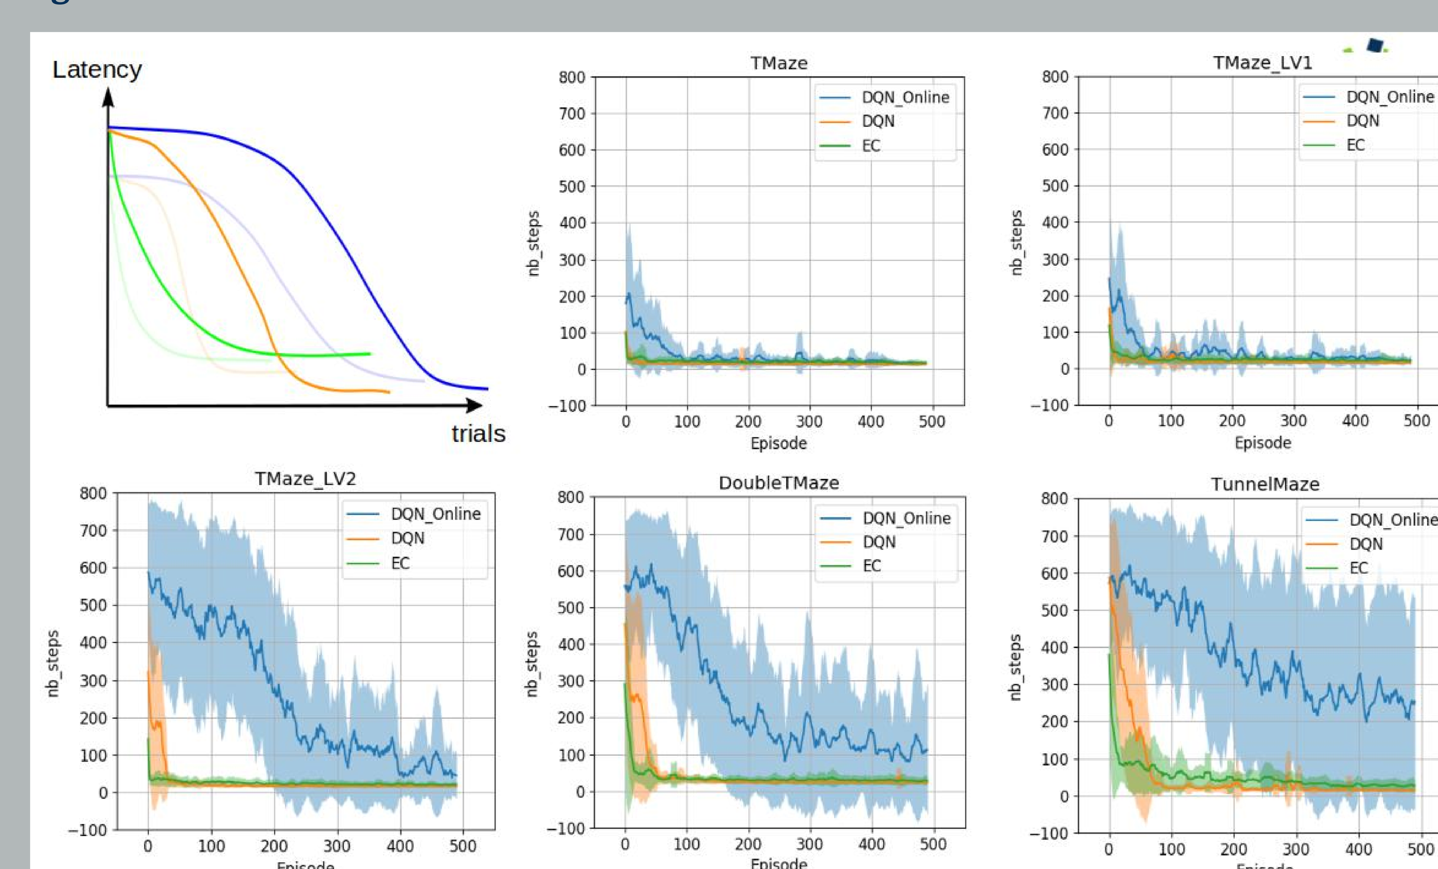

The plots show the learning curve for each agent (colored lines) inside each environment, as well as our theoretical expectations. The Y-axis indicates the number of steps the agent takes to find the reward and the X-axis the number of training episodes.

Conclusion

- The learning curves of the three agents confirm our expectations of the three learning paradigms
- Online deep Q learning is very sensitive to task complexity
- In simple tasks like the T-Maze, the convergence speed of DQN and episodic control are comparable
- Episodic control learns faster initially in complex tasks, but is more prone to get stuck in a sub-optimal solution, when there are multiple solutions and the globally optimal one is harder to find.

What's next...

1. Quantify the task complexity/statistics more comprehensively
2. Incorporate a generative episodic memory module
3. Improve the semantic system
4. Use the model to explain neural activity

Aim

Built a framework that can explain and unify experimental results from multiple hippocampally dependent tasks

Acknowledgements

This work has been supported by grant no. 419037518 from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the DFG Research Unit "Constructing Scenarios of the Past" (FOR 2812).

References

[1] Blundell, C., Uria, B., Pritzel, A., Li, Y., Ruderman, A., Leibo, J., Rae, J., Wierstra, D., and Hassabis, D.. (2016) *Model-free episodic control.* arXiv preprint arXiv:1606.04460
[2] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A., Veness, J., Bellemare, M. G., Graves, ..., Hassabis, D.. (2015) *Human-level control through deep reinforcement learning.* Nature, 518(7540):529–533

FOR 2812 CONSTRUCTING SCENARIOS OF THE PAST

DFG Deutsche Forschungsgemeinschaft

RUHR UNIVERSITÄT BOCHUM RUB