Meta-Learned World Models for Rapid Environmental Adaptation: A Hippocampal-Inspired Approach to Model-Based Reinforcement Learning

Rishabh Mallik, Emre Neftci, and John Paul Strachan

Peter Grünberg Institute, Forschungszentrum Jülich, Germany

Biological agents exhibit remarkable flexibility when navigating novel environments, rapidly constructing spatial representations through hippocampal cognitive maps while making episodic memory by few-shot learning. Inspired by this, we propose a meta-learned world model framework for model-based reinforcement learning that enables rapid adaptation across diverse environments. While model-based RL traditionally suffers from distribution shift when environment dynamics change, our approach meta-trains a recurrent world model to learn an adaptive prior over environment dynamics, analogous to how the hippocampus maintains a flexible scaffold for encoding novel spatial experiences.

Our framework draws on recent neuroscientific evidence showing that the hippocampus functions as a meta-learning module, where grid cell representations provide a structured basis for one-shot spatial memory encoding and episodic retrieval. We implement this through a two-timescale learning architecture: an outer loop that meta-learns world model parameters across a distribution of environments (analogous to slow synaptic consolidation), and an inner loop where the model rapidly adapts to local dynamics through context-dependent recurrent states (analogous to fast activity-based learning in hippocampal circuits). Unlike prior meta-RL approaches that meta-learn policies directly, our method meta-learns the environment model itself, enabling few-shot transfer of predictive dynamics while maintaining explicit model-based planning capabilities.

The key innovation lies in treating environmental adaptation as an episodic memory problem: the meta-learned world model maintains a library of compressed environment representations that can be rapidly retrieved and composed when encountering novel but structurally similar environments. This is implemented through a memory-augmented recurrent dynamics model trained with gradient-based meta-learning (MAML), where recent environment transitions serve as episodic context for online model adaptation. Crucially, this approach mirrors the hippocampal mechanism of pattern separation and completion, where similar spatial contexts activate overlapping neural ensembles that support rapid generalization.

[1] A. Nagabandi et al., Int. Conf. Learn. Represent., 2019.

[2] Z. Rimon et al., Int. Conf. Learn. Represent., 2024.

[3] I. Clavera et al., Proc. Mach. Learn. Res., 87, 2477–2486, 2018.

[4] J. X. Wang et al., Nat. Neurosci., 21, 860–868, 2018.

[5] L. Ambrogioni and H. F. Ólafsdóttir, Trends Cogn. Sci., 27(8), 702–714, 2023.

[6] K. L. Stachenfeld et al., Nat. Neurosci., 20, 1643–1653, 2017.

[7] J. L. S. Bellmund et al., eLife, 5, e16534, 2016.

[8] C. Finn et al., Proc. 34th Int. Conf. Mach. Learn., 1126–1135, 2017.