



Meta-learning three-factor learning rules for structured credit assignment in recurrent neural networks

Dimitra Maoutsa 1, 2

1. Technical University of Berlin, Berlin, DE

2. Independent, DE

Biological neural networks learn to execute complex tasks using feedback that is both sparse and delayed, yet synaptic plasticity operates locally both in space and time. How can delayed, global reinforcement signals drive the precise synaptic updates required for successful task execution? Previous work tackling this problem has considered either continuous error signals that guide learning throughout the trial period, or hand-crafted local updates. Yet, the space of plasticity rules capable of supporting learning from delayed reinforcement, and how these rules relate to task structure and resulting representations remains largely unexplored.

Here, we introduce a meta-learning framework for discovering biologically plausible three-factor learning rules for structured credit assignment in recurrent neural networks. Each rule couples a reward prediction error signal to synapse-specific eligibility traces that accumulate nonlinear interactions of pre- and post-synaptic activity over task execution. We optimise the trace coefficients in a meta-optimisation loop that employs a REINFORCE-like gradient approximation for the plasticity parameter updates, avoiding memory intensive backpropagation-through-time over the entire learning period, while maintaining accurate gradient estimates. We systematically explore how resulting representations depend on the employed form of discovered plasticity rules, and how these rules relate to task parameters. Our work offers insights into biologically grounded mechanisms for learning in recurrent circuits, showing how local plasticity interacts with reinforcement signals to give rise to behaviourally relevant representations.

NeuroAI, biologically plausible learning, plasticity rule inference, RNNs