Automatic Learning Rule Discovery for Neuromorphic Hardware

F. Assmuth, J. Lohoff, E. Neftci

Forschungszentrum Jülich, RWTH Aachen

Gradient-based optimization lies at the heart of modern Machine Learning (ML). Automatic Differentiation (AD) is the primary tool used to compute these gradients, usually with the backpropagation algorithm due to its simplicity and heuristically good fit. However the enormous compute, memory and communication costs of backprop and AD overall prevents non-von Neumann models from learning on-chip. The problem is thus often offline and inherently not bio-plausible on von Neumann architectures. Experimentally observed learning rules, such as Hebbian three-factor rules seem to provide an alternative, but lack the same empirical performance as gradient-based learning. Neuro-inspired gradient-based approximations, such as e-prop, have successfully connected both domains through retaining bio-plausibility and hardware efficiency, while still converging well – a zero cost approximation. However, the issue with these algorithms is the manual derivation and implementation of the Jacobian, as well as the sparse approximations, hindering them at scale.

$$\frac{\partial L}{\partial w} = \underbrace{\frac{\partial L}{\partial s}}_{\text{global}} \cdot \underbrace{\frac{\partial s}{\partial u}}_{\text{post}} \cdot \underbrace{\frac{\partial u}{\partial w}}_{\text{pre}}$$

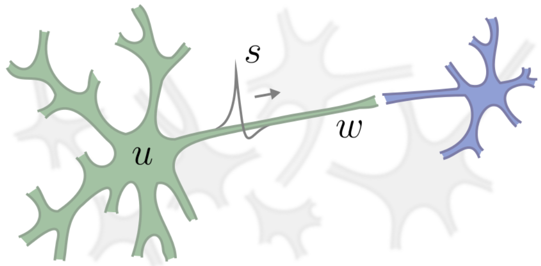

Where can approximations be exploited here?

- hardware utilization
- efficiency in bandwidth, memory, and compute

at zero cost!

Architecture specific

von Neumann	non-von Neumann
long range learning	on-chip learning

Figure 1: The neuro-inspired three factor rule can be viewed as the accumulation of Jacobians (factors): the global (gray), pre-synaptic (green) and post-synaptic (blue). Where L is the loss, s is the spike, u is the action potential, and w is the weight. The global factor is dependent on the choice of loss function, the post-synaptic factor is best taken as the surrogate gradient, and finally the pre-synaptic factor is the portion with which we can apply the learning rule discovery.

In previous works we demonstrated how we can improve basic AD modes such as backprop through a more optimized Jacobian accumulation order via Reinforcement Learning (RL) and apply it on ML and neuro-inspired models. The goal of this work is to extend the RL search problem to a data dependent approximation via systematic sparsification of Jacobians. Due to the nature of RL we are able to decrease the search to specific hardware restrictions, such as the memory footprint in tensor contractions. Furthermore we demonstrate minimal trade-offs in performance versus accuracy, testing on varying levels of synthetically correlated datasets as well as with pMNIST. We thus aim to demonstrate the ability to generalize the discovery of new complex learning rules, in a problem and hardware aware manner.

[1] Lohoff et al. *A truly sparse implementation of gradient-based plasticity*. NICE, 2025.

[2] Lohoff et al. *Optimizing Automatic Differentiation with Deep Reinforcement Learning*. NeurIPS, 2024.

[3] Bellec et al. *A solution to the learning dilemma for recurrent networks of spiking neurons*. Nature communications, 11, 2020.