Layer-Local Supervised Training of Spiking Neural Networks via Noise Perturbation

Mehrzad Karamimanesh, Marcel van Gerven, Mahyar Shahsavari

Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, the Netherlands

Spiking neural networks (SNNs) are a biologically-inspired computation model in which neurons communicate through discrete spike events, making them well-suited for energy-efficient neuromorphic hardware [1]. Despite this advantage, training SNNs remains challenging. The dominant supervised training approach combines backpropagation-through-time (BPTT) with surrogate gradient methods to handle the non-differentiability of spike functions [2]. This requires unrolling network dynamics across time steps and storing the full temporal activation history, resulting in memory costs that grow with both network depth and simulation duration. More critically, BPTT requires a global backward pass that stalls the computational pipeline, making it incompatible with online learning from continuous data streams [3].

Local learning rules offer a hardware-compatible alternative by restricting each synapse's update to locally available signals. However, most local rules for SNNs either rely on unsupervised Hebbian objectives that cannot directly optimize a task-specific loss, or require constructing target spike trains that are difficult to define for hidden layers. The challenge of assigning credit to hidden neurons without a backward pass remains a central open question in neuromorphic learning [3].

This work proposes a supervised online local learning algorithm for SNNs that extends activity-based node perturbation (ANP) [4] to spiking dynamics. Rather than injecting noise into continuous activations, small stochastic perturbations are applied to the membrane potential of spiking neurons, and the resulting change in a scalar task loss is used to estimate a local update direction without any backward computation. The algorithm operates in an online fashion, updating synaptic weights after each input presentation using only a forward pass and a scalar reward signal, with no temporal unrolling and no global gradient transport.

The local credit assignment problem is addressed through a layerwise critic structure in which each hidden layer is augmented with a lightweight linear readout trained to predict the task label from local spike rates. These critics provide the scalar goodness signal for perturbation-based updates without requiring information from deeper layers. A confidence-gated agreement mechanism weakly couples adjacent critics through local probability alignment, allowing task-relevant information to propagate forward without violating the locality constraint. To prevent simultaneous update resonance across layers during hardware execution, agreement updates are applied on an alternating schedule. The proposed method is designed for evaluation on neuromorphic benchmarks including N-MNIST and the Spiking Heidelberg Digits (SHD) dataset [5], targeting competitive accuracy under strict locality and online execution constraints.

[1] Karamimanesh et al., Neural Netw., 186, 107256, 2025.

[2] Eshraghian et al., Proc. IEEE, 111(9), 1016–1054, 2023.

[3] Zenke and Neftci, Proc. IEEE, 109(5), 935–950, 2021.

[4] Dalm et al., arXiv preprint arXiv:2310.00965, 2023.

[5] Cramer et al., IEEE Trans. Neural Netw. Learn. Syst., 33(7), 2744–2757, 2022.