

Scalable event-driven adaptation of supervised learning for spiking neural networks

A. Korcsak-Gorzo^(1,2), J. A. Espinoza Valverde⁽³⁾, J. Stapmanns⁽⁴⁾,
H. E. Plesser^(5,1,6), D. Dahmen⁽¹⁾, M. Bolten⁽³⁾, S. J. van Albada^(1,7), and M. Diesmann^(1,2,8,9)

⁽¹⁾ Institute for Advanced Simulation 6 (IAS-6), Jülich Research Centre, Jülich, Germany, ⁽²⁾ Department of Physics, Faculty 1, RWTH Aachen University, Aachen, Germany, ⁽³⁾ Department of Mathematics and Science, University of Wuppertal, Wuppertal, Germany, ⁽⁴⁾ Department of Physiology, University of Bern, Bern, Switzerland, ⁽⁵⁾ Department of Data Science, Faculty of Science and Technology, Norwegian University of Life Sciences, Ås, Norway, ⁽⁶⁾ Käte Hamburger Kolleg, RWTH Aachen University, Aachen, Germany, ⁽⁷⁾ Institute of Zoology, University of Cologne, Cologne, Germany, ⁽⁸⁾ JARA-Institute Brain Structure-Function Relationships (INM-10), Jülich Research Centre, Jülich, Germany, ⁽⁹⁾ Department of Psychiatry, Psychotherapy and Psychosomatics, Medical Faculty, RWTH Aachen University, Aachen, Germany

Biologically plausible learning unfolds over long timescales in large networks, requiring scalable and efficient algorithms to enable meaningful comparisons with neurobiological data. NEST [1] is an open-source simulator optimized for distributed large-scale spiking neural network simulations.

In our work [2], we port the supervised variant of eligibility propagation (e-prop) [3], a three-factor learning algorithm for recurrent spiking neural networks, from TensorFlow [4] to NEST by converting the time-driven weight update scheme into an event-driven scheme. This approach faithfully reproduces the original results on regression and classification tasks. In our implementation, the three factors are stored until a spike occurs, while ineffective history is discarded to reduce computational overhead. Additionally, memory usage is reduced by automatically triggered intermediate gradient computations for infrequently firing presynaptic neurons.

We further improve the model’s biological realism in several respects. Removing fixed update intervals eliminates globally synchronous updates and enables flexible processing of dynamic inputs. Inspired by truncated algorithms [5], each spike triggers a weight update, while neuronal and synaptic dynamics evolve continuously without resets and learning windows remain decoupled from global timing. Mean-squared error for classification [6] avoids communication in the output layer. The implementation also supports smoother surrogate gradients [7] and enforces Dale’s law by constraining synaptic weights. It maintains performance and efficiency on the neuromorphic MNIST dataset [8] and scales to millions of neurons.

Many biologically plausible learning rules are developed in time-driven frameworks optimized for continuous-valued communication. Our work provides a blueprint for translating such rules into event-driven implementations, which are particularly efficient in biologically realistic settings with sparse communication.

- [1] D. Terhorst et al., Zenodo, NEST 3.9, 2025.
- [2] A. Korcsak-Gorzo et al., arXiv, 2511.21674, 2025.
- [3] G. Bellec et al., Nat. Commun., 11 (1), 3625, 2020.
- [4] M. Abadi et al., arXiv, 1603.04467, 2016.
- [5] B. Yin et al., Nat. Mach. Intell., 5 (5), 518–527, 2023.
- [6] L. Hui and M. Belkin, arXiv, 2006.07322, 2021.
- [7] E. O. Neftci et al., IEEE Signal Process. Mag., 36 (6), 51–63, 2019.
- [8] G. Orchard et al., Front. Neurosci., 9, 437, 2015.