

Beyond Backprop: Teaching AI to Learn Like the Brain

Marco P. E. Apolinario

July 16, 2026

Abstract

Deep neural networks today are predominantly trained using the backpropagation algorithm. Although backpropagation has driven remarkable successes in artificial intelligence, it requires global error propagation and global weight synchronization, rendering it highly energy-intensive and biologically implausible. This paper explores the development of biologically inspired local learning rules that replace global backpropagation with mechanisms relying only on locally accessible signals. We discuss three key frameworks: S-TLLR for temporal credit assignment, LLS for spatial layer-wise learning via neural synchronization, and TESS, which unifies spatial and temporal locality into a single scalable rule for spiking neural networks. Together, these frameworks pave the way for energy-efficient, lifelong on-device learning.

1 Overview

Deep neural networks today are powered by the backpropagation algorithm — a mathematical engine that computes global gradients by transmitting information backward through every layer and time step. Despite its success, this mechanism is neither efficient nor biologically plausible. It requires every neuron to know the state of all others, making learning energy-hungry and unsuitable for on-device intelligence. The human brain, in contrast, learns **locally**: each synapse adjusts its strength using only signals it can access — pre- and post-synaptic activity and modulatory feedback.

My research explores a simple yet profound question:

Can we teach artificial systems to learn the way the brain does — locally, efficiently, and continuously?

This question guides my exploration of **local learning rules** — biologically inspired algorithms that replace global gradient propagation with mechanisms relying only on local information. The goal is to bring the power of deep learning closer to the remarkable efficiency and adaptability of the brain, enabling AI that can learn autonomously and on-device.

The first step toward this goal was to address **temporal locality** — assigning credit in time without replaying an entire sequence of neural activations. In **S-TLLR (Spike-Timing-Dependent Local Learning Rule)** [1], we introduced a temporal local learning mechanism inspired by *spike-timing-dependent plasticity (STDP)*. Instead of backpropagating errors through time as in BPTT, S-TLLR computes instantaneous updates from causal and non-causal spike correlations modulated by a local learning signal. This design allows training to proceed online, using constant memory and computation, while matching backpropagation accuracy within a few percentage points. By relying only on locally accessible quantities, S-TLLR brings time-local

learning to deep spiking models, opening a path toward scalable, energy-efficient on-device adaptation.

Yet, while S-TLLR achieved temporal locality, it still relied on global feedback across layers. To achieve **spatial locality**, we developed **LLS (Local Learning via Synchronization)** [2], where each layer learns through coordinated oscillatory activity rather than backpropagated gradients. Inspired by synchronization phenomena observed in neural populations, LLS allows deep networks to train layer-wise using fixed local feedback, achieving competitive accuracy on benchmarks such as CIFAR-100 and TinyImageNet. This revealed that effective representation learning can arise from *self-organized synchronization* rather than explicit error signals.

Building on these principles, **TESS (Temporally and Spatially Local Learning Rule)** [3] unifies both time and space into a fully local learning framework. TESS combines STDP-like eligibility traces for temporal credit assignment with layer-wise synchronization mechanisms for spatial coordination, allowing each neuron to update its synapses using only locally available signals. Remarkably, TESS performs within 1.4 percentage points of backpropagation on event-based datasets such as CIFAR10-DVS and IBM DVS Gesture, while requiring over **200× fewer computations** and **up to 10× less memory**. This marks a step change in the design of neuromorphic algorithms — making on-device learning both practical and biologically grounded.

Together, these studies trace the evolution from global gradient-based optimization toward **biologically inspired local learning**. They show that the key properties of deep learning — credit assignment, generalization, and continual adaptation — can emerge from local interactions between neurons, without any need for global error propagation. By coupling these insights with hardware-software co-design, our research envisions a future where intelligent systems learn autonomously, efficiently, and continuously — moving AI **beyond backpropagation** and closer to the learning principles of the brain.

References

- [1] M. P. E. Apolinario and K. Roy, “S-TLLR: STDP-inspired Temporal Local Learning Rule for Spiking Neural Networks,” *Transactions on Machine Learning Research (TMLR)*, 2025.
- [2] M. P. E. Apolinario, A. Roy, and K. Roy, “LLS: Local Learning Rule for Deep Neural Networks Inspired by Neural Activity Synchronization,” in *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2025.
- [3] M. P. E. Apolinario, K. Roy, and C. Frenkel, “TESS: A Scalable Temporally and Spatially Local Learning Rule for Spiking Neural Networks,” in *International Joint Conference on Neural Networks (IJCNN)*, 2025.