Dendroprop: Dendritic plateau potentials link somatic and dendritic inputs over long timescales using a surrogate-gradient-based learning rule

T. Bax(1), P. Jain (2), C.J. Do Nascimento(2), J. Senk(2,3) and P. Nieters(1)

(1) Institute of Cognitive Science, Osnabrück University, (2) Sussex AI, School of Engineering and Informatics, University of Sussex (3) Institute for Advanced Simulation (IAS-6), Jülich Research Centre,

The simple, single-compartment Leaky Integrate-and-Fire (LIF) neuron is the standard model of spike-based computation in neuroscience and machine learning. This is despite overwhelming evidence that computation in pyramidal neurons is strongly driven by active dendritic processes that respond to inputs locally in complex dendritic trees [1]. This gap persists because the field lacks clear principles for dendritic contributions to computation. Here, we investigate NMDA-mediated dendritic plateau potentials and hypothesize that their key contribution is the ability to bind temporally dispersed spike patterns robustly over long time windows: When the dendritic compartment detects a transient spike pattern, it generates a plateau potential, resulting in a neural UP state. This state biases pattern recognition in the somatic compartment towards patterns occurring within one plateau duration, thereby binding dendritic and somatic patterns into one output response.

To test our hypothesis, we developed a two-compartment model consisting of a LIF soma coupled to a dendritic compartment that generates plateau potentials. The plateau acts on the soma by halving the firing threshold, modeling the elevated excitability of a neural UP state for the duration of a plateau—approximately one order of magnitude longer than a typical membrane response to an afferent spike.

We then derived a new surrogate gradient descent learning rule [2] for this model. The key property is that we assign credit to dendritic synapses based on whether somatic spikes occur during the plateau, regardless of their specific timing. Our model can thus learn to associate spike patterns up to a plateau length apart, enabling effective learning of long-range temporal dependencies.

Finally, we evaluated our model against a parameter-matched LIF network on two benchmark datasets. On the Neuromorphic-MNIST dataset [3], which can be solved without temporal information, our model matched the LIF network performance (both 98% accuracy). On the Spiking Heidelberg Digits [4], which requires temporal information, our model improved significantly over the LIF network (70% vs. 48% accuracy), demonstrating that plateaus significantly aid temporal processing. In both cases, our model learned various temporally dispersed pattern pairs as features, supporting our hypothesis that plateaus bind spike patterns for recognition.

In sum, we showed that dendritic plateaus turn neurons into multi-timescale pattern detectors. This general computational function aligns with specific capabilities assigned to dendrites, such as memory traces [5] or sequence processing [6]. The clear advantages of our plateau-then-spike model over LIF neurons strongly advocate for the integration of dendritic computation into neuron models from both practical and theoretical perspectives.

[1] M. London & M. Häusser, Annu. Rev. Neurosci., 28 (1), 503-532, 2005.

[2] E. O. Neftci et al., IEEE Signal Process. Mag., 36 (6), 51-63, 2019.

[3] G. Orchard et al., Front. Neurosci., 9, 437, 2015.

[4] B. Cramer et al., IEEE Trans. Neural Netw. Learn. Syst., 33 (7), 2744-2757, 2020.

[5] A. Quaresima et al., J. Physiol., 601 (15), 3265-3295, 2023.

[6] J. Leugering et al., Front. Cogn., 2, 1044216, 2023.