

Learning Precise Spike Timings with Eligibility Traces

Manuel Traub¹, Martin V. Butz¹, R. Harald Baayen², and Sebastian Otte¹

¹ University of Tübingen – Neuro-Cognitive Modeling Group,
Sand 14, 72076 Tübingen, Germany

² University of Tübingen – Quantitative Linguistics,
Wilhelmstr. 19, 72074 Tübingen, Germany

Abstract. Recent research in the field of spiking neural networks (SNNs) has shown that recurrent variants of SNNs, namely *long short-term SNNs* (LSNNs), can be trained via error gradients just as effective as LSTMs. The underlying learning method (e-prop) is based on a formalization of eligibility traces applied to *leaky integrate and fire* (LIF) neurons. Here, we show that the proposed approach cannot fully unfold *spike timing dependent plasticity* (STDP). As a consequence, this limits in principle the inherent advantage of SNNs, that is, the potential to develop codes that rely on precise relative spike timings. We show that STDP-aware synaptic gradients naturally emerge within the eligibility equations of e-prop when derived for a slightly more complex spiking neuron model, here at the example of the Izhikevich model. We also present a simple extension of the LIF model that provides similar gradients. In a simple experiment we demonstrate that the STDP-aware LIF neurons can learn precise spike timings from an e-prop-based gradient signal.

Keywords: eligibility traces · recurrent neural networks · backpropagation through time · spike timing dependent plasticity

1 Introduction

Spike Timing Dependent Plasticity (STDP) is assumed to be a fundamental learning principle in the brain [5]. It is considered a prerequisite for developing temporal codes in which precise relative spike timings are key, thus going beyond plain rate coding. Accordingly, STDP is based on temporal correlations between presynaptic and postsynaptic neural activities. Indeed, it has been shown that such a Hebbian-like form of synaptic plasticity is at play in the visual cortex of primates [8].

Recently, (non-supervised) STDP based learning rules were successfully applied to train deep (non-recurrent) convolutional SNNs for image recognition [10,13]. In terms of effectiveness, though, supervised back-propagation-like approaches seem to be more promising [11], even though spiking neurons are not differentiable per se, such that an error gradient signal can only be approximated.

Recently, it has been shown that even *back-propagation through time* (BPTT) [14] can be applied for training recurrent SNNs [1]. Bellec et al. demonstrated

that SNNs, specifically a variant that is called *long short-term SNN* (LSNN), for the first time can reach the performance of the well-known LSTM [7]. Moreover, the mathematical approach of learning in SNNs has led to the derivation of a biologically plausible learning rule called *e-prop*, which approximates BPTT [2] by means of a formalization of *eligibility traces*.

Bellec et al. established a link between e-prop and (biological) synaptic plasticity [4,3]. It appears, however, that STDP can only fully arise within eligibility trace-based learning, when the neuron model provides a negated gradient signal in the case when a presynaptic spike arrives too late, i.e., shortly after a postsynaptic spike. This does not happen in the LIF model proposed in [3], cf. Section 4 for further details.

The contributions of this paper are as follows. We show that STDP emerges within e-prop when it is derived using a more complex neuron model, which adequately incorporates a refractory period. This is exemplarily shown for the Izhikevich model [9]. Moreover, we present an adjustment of the basic LIF model used in LSNNs in order to produce the same STDP behavior within the e-prop framework.

2 Background

Experimental data suggests that the brain solves the *temporal credit assignment problem* by combining local eligibility traces, which maintain information about individual synapses' activation histories, with neuromodulator-based reward signals [6]. The e-prop algorithm [2] adapts this principle by factorizing the error gradients from BPTT into a sum of products between local eligibility traces and online learning signals.

$$\begin{aligned} \frac{dE}{dw_{i,j}} &= \sum_t L_j^t e_{i,j}^t \\ &= \sum_t \frac{dE}{dz_j^t} e_{i,j}^t \end{aligned} \quad (1)$$

Here, the learning signal (L_j^t) represents the global error information of the post-synaptic spike, whereas the eligibility trace ($e_{i,j}^t$) captures the local information available at the synapse. z_j^t refers to the (spiking) output of neuron j at time step t .

The eligibility trace is furthermore a product of pre- and postsynaptic information, but does not include any error gradient information:

$$e_{i,j}^t = \frac{\partial z_j^t}{\partial \mathbf{s}_j^t} \boldsymbol{\epsilon}_{i,j}^t \quad (2)$$

Specifically, it is the product of a *pseudo derivative*, replacing the non-existing derivative of the spiking function $\partial z_j^t / \partial \mathbf{s}_j^t$ and the presynaptic activity flow accumulated within an eligibility vector $\boldsymbol{\epsilon}_{i,j}^t$. The latter is computed forward through

time:

$$\epsilon_{i,j}^t = \epsilon_{i,j}^{t-1} \frac{\partial \mathbf{s}_j^t}{\partial \mathbf{s}_j^{t-1}} + \frac{\partial \mathbf{s}_j^t}{\partial w_{i,j}} \quad (3)$$

where \mathbf{s}_j^t refers to the current state of a neuron containing, for instance (depending on the model), its action potential and possibly further adaptive parameters.

3 STPD with Izhikevich Neurons

The Izhikevich neuron [9] is a precise, but computationally cheap model of a biological neuron that uses two parameterizable differential equations. It is particularly more complex than the simple LIF model, but, more importantly, explicitly models the refractory period of the neuron.

3.1 Izhikevich Model and Eligibility Trace

The dynamics of the Izhikevich neuron are described with the following differential equations.

$$v' = 0.04v^2 + 5v + 140 - u + I \quad (4)$$

$$u' = 0.004v - 0.02u \quad (5)$$

Here v is the membrane voltage and u is a recovery variable, which controls the refractory period of the Izhikevich neuron. I is the current input to the neuron.

Once the membrane voltage crosses 30mV, a spike is emitted and v and u are reset as described in Algorithm 1.

In order to derive an eligibility trace for Izhikevich neurons, the equations (4) and (5) have to be modeled in discrete time steps, and also the reset has to be modeled within the new equations. To accomplish this built-in reset, the variables \tilde{v}_j^t and \tilde{u}_j^t are introduced, replacing v and u in the standard equations.

$$\tilde{v}_j^t = v_j^t - (v_j^t + 65)z_j^t \quad (6)$$

$$\tilde{u}_j^t = u_j^t + 2z_j^t \quad (7)$$

Here the binary variable z_j^t , resets \tilde{v}_j^t and \tilde{u}_j^t after a spike of neuron j at time step t . Now v and u can be computed in discrete time steps using Euler integration with a constant step size of δt .

$$v_j^{t+1} = \tilde{v}_j^t + \delta t(0.04(\tilde{v}_j^t)^2 + 5\tilde{v}_j^t + 140 - \tilde{u}_j^t + I_j^t) \quad (8)$$

Algorithm 1 Izhikevich neuron reset

```

if  $v < 30mV$  then
   $v \leftarrow -65mV$ 
   $u \leftarrow u + 2$ 
end if

```

$$u_j^{t+1} = \tilde{u}_j^t + \delta t(0.004\tilde{v}_j^t - 0.02\tilde{u}_j^t) \quad (9)$$

The hidden state of Izhikevich neurons is then defined as a two-dimensional vector containing v_j^t and u_j^t .

$$\mathbf{s}_j^t = \begin{pmatrix} v_j^t \\ u_j^t \end{pmatrix} \quad (10)$$

To finally derive the eligibility vector (3), the derivative of the next hidden state \mathbf{s}_j^{t+1} by the current state \mathbf{s}_j^t has to be computed. This state derivative can be expressed in the following form of a 2x2 matrix.

$$\frac{\partial \mathbf{s}_j^{t+1}}{\partial \mathbf{s}_j^t} = \begin{pmatrix} \frac{\partial v_j^{t+1}}{\partial v_j^t} & \frac{\partial v_j^{t+1}}{\partial u_j^t} \\ \frac{\partial u_j^{t+1}}{\partial v_j^t} & \frac{\partial u_j^{t+1}}{\partial u_j^t} \end{pmatrix} \quad (11)$$

The partial derivatives can be further simplified by taking into account that z_j^t is a binary variable.

$$\begin{aligned} \frac{\partial v_j^{t+1}}{\partial v_j^t} &= 1 - z_j^t \\ &\quad + 0.08\delta t(v_j^t - (v_j^t + 65)z_j^t)(1 - z_j^t) \\ &\quad + 5\delta t(1 - z_j^t) \end{aligned} \quad (12)$$

$$\begin{aligned} &= 1 - z_j^t + 0.08\delta t v_j^t(1 - z_j^t) \\ &\quad + 5\delta t(1 - z_j^t) \end{aligned}$$

$$= (1 - z_j^t)(1 + (0.08v_j^t + 5)\delta t)$$

$$\frac{\partial v_j^{t+1}}{\partial u_j^t} = -\delta t \quad (13)$$

$$\frac{\partial u_j^{t+1}}{\partial v_j^t} = 0.004\delta t(1 - z_j^t) \quad (14)$$

$$\frac{\partial u_j^{t+1}}{\partial u_j^t} = 1 - 0.02\delta t \quad (15)$$

Given this state derivative, the eligibility vector is computed in the following two-dimensional vector, which contains a voltage eligibility value in the first row and a refractory eligibility value in the second row:

$$\boldsymbol{\epsilon}_{i,j}^{t+1} = \frac{\partial \mathbf{s}_j^{t+1}}{\partial \mathbf{s}_j^t} \cdot \boldsymbol{\epsilon}_{i,j}^t + \frac{\partial \mathbf{s}_j^{t+1}}{\partial \theta_{ji}^{rec}} \quad (16)$$

where

$$\begin{aligned} \epsilon_{i,j,v}^{t+1} = & (1 - z_j^t)(1 + (0.08v_j^t + 5)\delta t)\epsilon_{i,j,v}^t \\ & - \delta t\epsilon_{i,j,u}^t + \delta tz_i^t \end{aligned} \quad (17)$$

$$\epsilon_{i,j,u}^{t+1} = 0.004\delta t(1 - z_j^t)\epsilon_{i,j,v}^t + (1 - 0.02\delta t)\epsilon_{i,j,u}^t \quad (18)$$

Here the recovery eligibility value $\epsilon_{i,j,u}^t$ approximates an exponential filter of the voltage eligibility.

The voltage eligibility itself is a function of the neuron’s action potential that also accumulates presynaptic spikes and gets dampened by its exponential average (by the recovery eligibility).

Another insight from the equations is that, whenever the postsynaptic neuron spikes, the voltage eligibility vector is reset to the negative recovery eligibility vector.

In order to derive the final eligibility trace, a pseudo-derivative for the Izhikevich neuron is defined as follows:

$$h_j^t := \gamma \exp\left(\frac{\min(v, 30) - 30}{30}\right) \quad (19)$$

where γ is a damping factor.

With this pseudo-derivative, the neuron spike z_j^t derived by the hidden state \mathbf{s}_j^t is then defined as:

$$\begin{pmatrix} z_j^t \\ v_j^t \\ z_j^t \\ u_j^t \end{pmatrix} \stackrel{\text{def}}{=} \begin{pmatrix} h_j^t \\ 0 \end{pmatrix} \quad (20)$$

The eligibility trace then simplifies to the pseudo-derivative times the voltage eligibility vector.

$$\begin{aligned} e_{i,j}^{t+1} &= \frac{\partial z_j^{t+1}}{\partial \mathbf{s}_j^{t+1}} \cdot \boldsymbol{\epsilon}_{i,j}^{t+1} \\ &= \begin{pmatrix} h_j^{t+1} & 0 \end{pmatrix} \begin{pmatrix} \epsilon_{i,j,v}^{t+1} \\ \epsilon_{i,j,u}^{t+1} \end{pmatrix} \\ &= h_j^{t+1} \epsilon_{i,j,v}^{t+1} \end{aligned} \quad (21)$$

3.2 Experimental Results

In the following evaluation, our goal was to inspect the evolution of the derived eligibility traces. Two Izhikevich neurons were weakly connected by a synapse and received random inputs to simulate behavior within a greater network. To specifically investigate the influence of the eligibility trace on the gradient, a constant positive learning signal was used. The gradient was calculated based on (1).




Fig. 1. Simulation of two connected Izhikevich neurons that exhibit a positively rewarded STDP behavior followed by a negatively rewarded one. As a result, a gradient computed with a constant positive learning signal increases during positively rewarded STDP and decreases during negatively rewarded STDP.

In the simulation shown in Figure 1, first an artificial strengthening STDP behavior is introduced by using an overall lower random input current to the output neuron. To ensure that an output spike (postsynaptic neuron) fires shortly after an input spike (presynaptic neuron) we steadily increased the random input current for the output neuron after the input neuron spiked. In the second part of the simulation this behavior is reversed (presynaptic spike before postsynaptic spike) in order to produce a weakening STDP behavior, as shown in Algorithm 2. Here $U(\alpha, \beta)$ represents a random variable drawn uniform from the interval (α, β) and t_{z_i}, t_{z_o} are the spike times for the last input or output spike.

As can be seen in the diagram, the desired STDP behavior of the synapse is clearly reflected within the gradient, which for the first half of the simulation increases, and then decreases. The eligibility trace and the eligibility vector also reflect the introduced STDP behavior. They have only positive values during the STDP strengthening phase, but they also take on negative values during the STDP weakening phase.

Algorithm 2 Izhikevich e-prop simulation

```

if  $t < T \cdot 0.45$  then
   $I_i^t \leftarrow U(1, 15)$ 
   $I_o^t \leftarrow U(1, 5)$ 
  if  $t_{z_i} > t_{z_o}$  then
     $I_o^t \leftarrow U(0, 1)(t - t_{z_i})$ 
  end if
else
   $I_i^t \leftarrow U(1, 5)$ 
   $I_o^t \leftarrow U(1, 15)$ 
  if  $t_{z_o} > t_{z_i}$  then
     $I_i^t \leftarrow U(0, 1)(t - t_{z_o})$ 
  end if
end if

```

4 STDP with LIF Neurons

While our experimental results confirm the desired general tendency to reflect full STDP behavior in gradients based on the eligibility traces of Izhikevich neurons, plain LIF neurons do not show this behavior. We now detail the reason for this lack and introduce negative eligibilities to induce effective connection weakening.

4.1 LIF Model

Considering the definition of the dynamics of the LIF neuron

$$v_j^{t+1} = \alpha v_j^t + I_j^t - z_j^t v_{thr} \quad (22)$$

Equation (22) shows the action potential dynamics of a LIF neuron in discrete timesteps, as defined in [2]. The neuron integrates over the weighted sum of incoming spikes I_j^t ; $\alpha < 1$ controls the leakage. Once a LIF neuron spikes, its action potential is reset by subtracting the value of the spike threshold.

A neuron spike is modeled by the Heaviside step function. The neuron is prohibited from spiking during a fixed refractory period after the last spike.

$$z_j^t = \begin{cases} 0, & \text{if } t - t_{z_j} < \delta t_{ref} \\ H(v_j^t - v_{thr}), & \text{otherwise} \end{cases} \quad (23)$$

Here (23) t_{z_j} represents the most recent spike time of neuron j and δt_{ref} denotes the length of the refractory period.

Seeing that also the *pseudo derivative* (24) is set to zero during this refractory period:

$$h_j^t := \begin{cases} 0, & \text{if } t - t_{z_j} < \delta t_{ref} \\ \gamma \max(0, 1 - |\frac{v_j^t - v_{thr}}{v_{thr}}|), & \text{otherwise} \end{cases}, \quad (24)$$

no weakening STDP based gradients can unfold in the e-prop equations using standard LIF neurons.

4.2 STDP-LIF Model and Eligibility Trace

Whereas an STDP influenced gradient cannot be observed with the standard equations for LIF or adaptive LIF (ALIF) neurons within an LSNN [1], by slightly modifying the original LIF formulation from (22), a clear STDP based eligibility trace emerges.

In order to compute such an eligibility trace reflecting the STDP behavior of the synapse connecting two LIF neurons, the LIF equation has to be slightly altered into, what we from now on will refer to, an STDP-LIF neuron:

$$v_j^{t+1} = \alpha v_j^t + I_j^t - z_j^t \alpha v_j^t - z_j^{t-\delta t_{ref}} \alpha v_j^t \quad (25)$$

Instead of using a soft reset at a fixed threshold, in (25) the STDP-LIF neuron is hard reset to zero whenever it spikes, and whenever its refractory period ends. Since the reset now does include the voltage v_j^t as a factor, the spike is now included in the hidden state derivative

$$\begin{aligned} \frac{\partial v_j^{t+1}}{\partial v_j^t} &= \alpha - z_j^t \alpha - \alpha z_j^{t-\delta t_{ref}} \\ &= \alpha(1 - z_j^t - z_j^{t-\delta t_{ref}}) \end{aligned} \quad (26)$$

and hence also in the computation of the eligibility trace

$$\begin{aligned} \epsilon_{i,j}^{t+1} &= \frac{\partial s_j^{t+1}}{\partial s_j^t} \cdot \epsilon_{i,j}^t + \frac{\partial s_j^{t+1}}{\partial \theta_{ji}^{rec}} \\ &= \alpha(1 - z_j^t - z_j^{t-\delta t_{ref}}) \epsilon_{i,j}^t + z_i^t \end{aligned} \quad (27)$$

$$\begin{aligned} e_{i,j}^{t+1} &= \frac{\partial z_j^{t+1}}{\partial s_j^{t+1}} \cdot \epsilon_{i,j}^{t+1} \\ &= h_j^{t+1} \epsilon_{i,j}^{t+1} \end{aligned} \quad (28)$$

As a result, the eligibility trace is reset after a spike and after the refractory period.

This reset behavior allows us now to create an eligibility trace that reflects the STDP behavior of its synapse by using a constant negative pseudo-derivative during the refractory period and otherwise leave the neuron dynamics unchanged.

$$h_j^t := \begin{cases} -\gamma, & \text{if } t - t_{z_j} < \delta t_{ref} \\ \gamma \max(0, 1 - |\frac{v_j^t - v_{thr}}{v_{thr}}|), & \text{otherwise} \end{cases} \quad (29)$$

As a result, any incoming spike during the refractory period produces a negative eligibility trace that persists for the time of the refractory period and has a negative influence on the gradient.




Fig. 2. Simulation of two connected LIF neurons that exhibit a positively rewarded STDP behavior followed by a negatively rewarded one. As a result, a gradient computed with a constant positive learning signal increases during positively rewarded STDP and decreases during negatively rewarded STDP.

4.3 Experimental Results

The STDP influenced gradient is shown in Figure 2, in which two connected STDP-LIF neurons are simulated similarly to the Izhikevich neuron simulation in Algorithm 2. Strengthening and weakening STDP events can be clearly identified in the eligibility trace, directly influencing the resulting gradient. Thus, by the above simple modifications of the LIF model, one can derive eligibility traces for STDP-LIF neurons that reflect the STDP behavior of the underlying synapse.

Since the gradient calculation using eligibility traces in combination with a back-propagated learning signal (1) is mathematically equivalent to normal BPTT, it follows that BPTT itself facilitates STDP behavior.

4.4 Learning Precise Timing

In a simple additional experiment we evaluated the performance of STDP-LIF neurons in comparison to standard LIF neurons by learning an LSNN with 16 hidden neurons, one input neuron, and one output neuron to approximate a function based on the timing of the input spike. In this experiment the LSNN receives Poisson distributed spikes from the input neuron with an average spike rate of 25 Hz. The supervised target signal for the leaky readout neuron is then

$$v_{target}^t = \frac{1}{1 + t - t_{in}} \quad (30)$$




Fig. 3. Spike timing test: STDP-LIF neurons manage to learn a simple function based on the inputs spike timing efficiently, while standard LIF neurons produce a larger error by approximating the function using a high firing rate. Graphs are averaged over 100 independent runs each.

In this equation, t_{in} is the spike time of the most recent input spike. A good approximation of (30) can be learned when the input spikes are just forwarded to the readout neuron, which can in turn approximate the shape of the function via its readout decay.

In the simulation, the spike threshold v_{thr} for each neuron was set to 0.5 and the LSN was trained over 1000 epochs using a batch size of 16 and an initial learning rate of 0.003, which is decayed every 100 epochs by multiplying it with 0.7. The network was trained with Adam [12] and no further regularization was used.

As can be seen in Figure 3, the STDP-LIF neurons quickly manage to suppress the hidden connections and only forward the input spikes to the readout neuron, while the LIF neurons on the other hand produce a larger error by approximating the function with a high firing rate.

5 Conclusion

In this paper we showed that STDP emerges within eligibility trace-based gradient signals in SNNs, given that the neuron model is sufficiently detailed. Specifically, it is crucial that the refractory period of postsynaptic neurons is taken into account. While this is not the case for regular LIF neurons, we demonstrated that by including eligibility traces derived from the well-known Izhikevich model, the gradient signal induces STDP behavior. The standard LIF model can also be refined such that it provides STDP. Equipped with STDP-aware gradient signals, we showed that learning precise spike timings becomes possible.

Seeing the recent advances in applying (recurrent) SNNs successfully, we will evaluate these mathematical insights very soon to harder benchmark problems like handwriting or speech recognition (TIMIT) and expect to achieve similar performances. The observed effect of reduced firing rates when using STDP-aware gradients could be particularly interesting for neuromorphic hardware implementations, in which sparse spiking behavior effectively benefits energy consumption.

References

1. Bellec, G., Salaj, D., Subramoney, A., Legenstein, R., Maass, W.: Long short-term memory and learning-to-learn in networks of spiking neurons. In: *Advances in Neural Information Processing Systems*. pp. 795–805 (2018)
2. Bellec, G., Scherr, F., Hajek, E., Salaj, D., Legenstein, R., Maass, W.: Biologically inspired alternatives to backpropagation through time for learning in recurrent neural nets. *arXiv preprint arXiv:1901.09049* (2019)
3. Bellec, G., Scherr, F., Hajek, E., Salaj, D., Subramoney, A., Legenstein, R., Maass, W.: Eligibility traces provide a data-inspired alternative to backpropagation through time (2019)
4. Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., Maass, W.: A solution to the learning dilemma for recurrent networks of spiking neurons. *bioRxiv* p. 738385 (2019)
5. Caporale, N., Dan, Y.: Spike timing-dependent plasticity: a hebbian learning rule. *Annu. Rev. Neurosci.* **31**, 25–46 (2008)
6. Gerstner, W., Lehmann, M., Liakoni, V., Corneil, D., Brea, J.: Eligibility traces and plasticity on behavioral time scales: experimental support of neohebbian three-factor learning rules. *Frontiers in neural circuits* **12** (2018)
7. Hochreiter, S., Schmidhuber, J.: Long Short-Term Memory. *Neural Comput.* **9**(8), 1735–1780 (Nov 1997). <https://doi.org/10.1162/neco.1997.9.8.1735>
8. Huang, S., Rozas, C., Trevino, M., Contreras, J., Yang, S., Song, L., Yoshioka, T., Lee, H.K., Kirkwood, A.: Associative hebbian synaptic plasticity in primate visual cortex. *Journal of Neuroscience* **34**(22), 7575–7579 (2014)
9. Izhikevich, E.M.: Simple model of spiking neurons. *IEEE Transactions on neural networks* **14**(6), 1569–1572 (2003)
10. Kheradpisheh, S.R., Ganjtabesh, M., Thorpe, S.J., Masquelier, T.: Stdp-based spiking deep convolutional neural networks for object recognition. *Neural Networks* **99**, 56–67 (2018)
11. Kheradpisheh, S.R., Masquelier, T.: S4nn: temporal backpropagation for spiking neural networks with one spike per neuron. *arXiv preprint arXiv:1910.09495* (2019)
12. Kingma, D.P., Ba, J.L.: Adam: A method for stochastic optimization. *3rd International Conference for Learning Representations* **abs/1412.6980** (2015)
13. Mozafari, M., Kheradpisheh, S.R., Masquelier, T., Nowzari-Dalini, A., Ganjtabesh, M.: First-spike-based visual categorization using reward-modulated stdp. *IEEE Transactions on Neural Networks and Learning Systems* (2018)
14. Werbos, P.: Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE* **78**(10), 1550–1560 (1990). <https://doi.org/10.1109/5.58337>