SAND2022-15707C

Sandia National Laboratories

Deriving STDP from Backpropagation

Nicholas Alonso, Felix Wang, Corinne Teeter

11/9/2022

SNUFA 2022

U.S. DEPARTMENT OF ENERGY
NNSA

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Overview

When using leaky integrate and fire neurons and assuming neurons are attempting to predict the precise timing of post-synaptic neurons:

The plasticity rules derived from e-prop look similar to standard STDP curves

These plasticity rules can learn sequences

We suggest that STDP not only optimizes networks for prediction but does so by preforming credit assignment though time via an algorithm akin to e-prop.

Backpropagation through time with Eligibility Traces

Eligibility propagation (e-prop) is equivalent to BPTT.

Unlike BPTT, e-prop does not store computational graph through time.

Instead, only stores local *eligibility traces*.

- E traces store part of the global loss gradient describing network dynamics through time.
- Allows for online training
- More bio-plausible
- Less memory

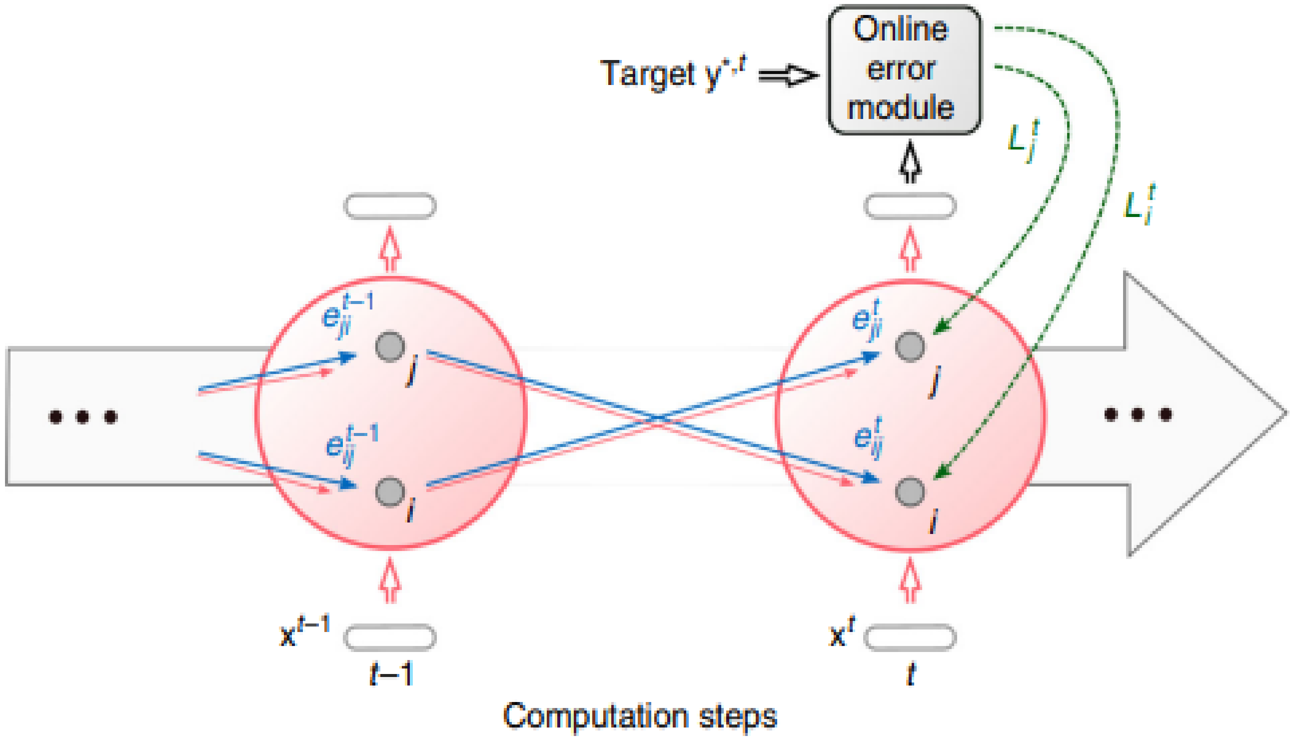

Bellec et al. , 2020

Eligibility trace (e_ij) stored for each synapse ij, and computed forward in time allowing for online updates (partially) equivalent to BPTT when combined with learning signal L.

Standard STDP Formulation

$$STDP = (1 - W_{ij})z_j^{t+1}a_i^t - W_{ij}z_i^t a_j^{t+1}$$

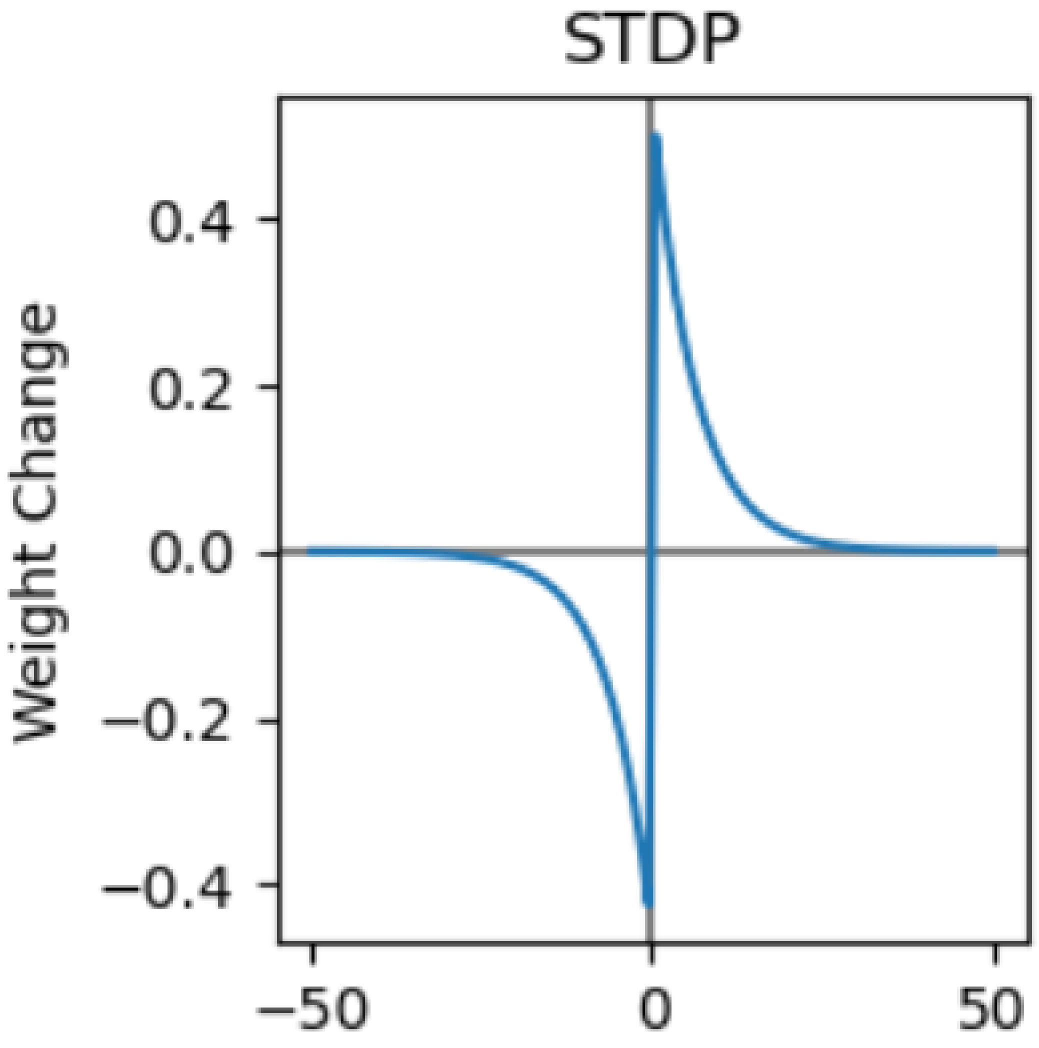

z_i: neuron i spike

a_i^t: neuron i trace: $z^t + \alpha a_i^{t-1}$

W_{ij}: weight from

Local Loss Function

Objective: Presynaptic neuron is trying to predict the spike of the post synaptic neuron

Loss

$$L = \frac{1}{2}\|e\|^2 = \frac{1}{2}\|z_j^t - \underbrace{W_{ij} z_i^{t-1}}_{prediction}\|^2$$

or

$$L = \frac{1}{2}\|e\|^2 = \frac{1}{2}\|z_j^t - h_j^t\|^2 = \frac{1}{2}\|z_j^t - \underbrace{(W_{ij} z_i^{t-1} + b^t)}_{prediction-like}\|^2$$

LIF Neuron Model w/ Reset

$$h^{t+1} = \alpha h^t + W z^t - (\alpha h^t + W z^t) z^t$$

L : Loss

i : index of presynaptic neuron

j : index of postsynaptic neuron

t : time

z : neuron spike (0 or 1)

h : membrane potential of neuron

Θ : Step/spike sampling function

W_{ij} : synaptic weight from i to j

e : error

b : placeholder for all of the terms in h that do not involve W

Derivation of pSTDP from Neuron Model 1

Eligibility Vector:

$$a^{t-1} = \sum_{t' \leq t} \frac{\partial h_j^t}{\partial h_j^{t-1}} \cdots \frac{\partial h_j^{t'+1}}{\partial h_j^{t'}} \frac{\partial h_j^{t'}}{\partial W_j} = \alpha a^{t-2} + z^{t-1}$$

Poisson Pseudo-Gradient:

$$\frac{\partial z^{t+1}}{\partial h^{t+1}} \approx 1 - a^{t+1}$$

p-STDP1

$$-\frac{\partial L}{\partial W_{ij}} = -\frac{\partial \frac{1}{2} \| \Theta(h_j^{t+1}) - h_j^{t+1} \|^2}{\partial h_j^{t+1}} \sum_{t' \leq t} \frac{\partial h_j^{t+1}}{\partial h_j^t} \cdots \frac{\partial h_j^{t'+1}}{\partial h_j^{t'}} \frac{\partial h_j^{t'}}{\partial W_j}$$

$$= -\frac{\partial L}{\partial e} \left(\frac{\partial e}{\partial z_j^{t+1}} \frac{\partial z_j^{t+1}}{\partial h_j^{t+1}} + \frac{\partial e}{\partial h_j^{t+1}} \right) a_i^t$$

$$\approx -(z_j^{t+1} - h_j^{t+1})(1 - a_j^{t+1} - 1) a_i^t$$

$$\ldots \text{(simplify)}$$

$$= \underbrace{z_j^{t+1} a_i^t - W_{ij} z_i^t a_j^{t+1}}_{STDP} - \underbrace{b a_j^{t+1} a_i^t}_{modulation}$$

t: timestep

z: neuron spikes

W: synaptic weight matrix

α: decay rate

h: membrane potential

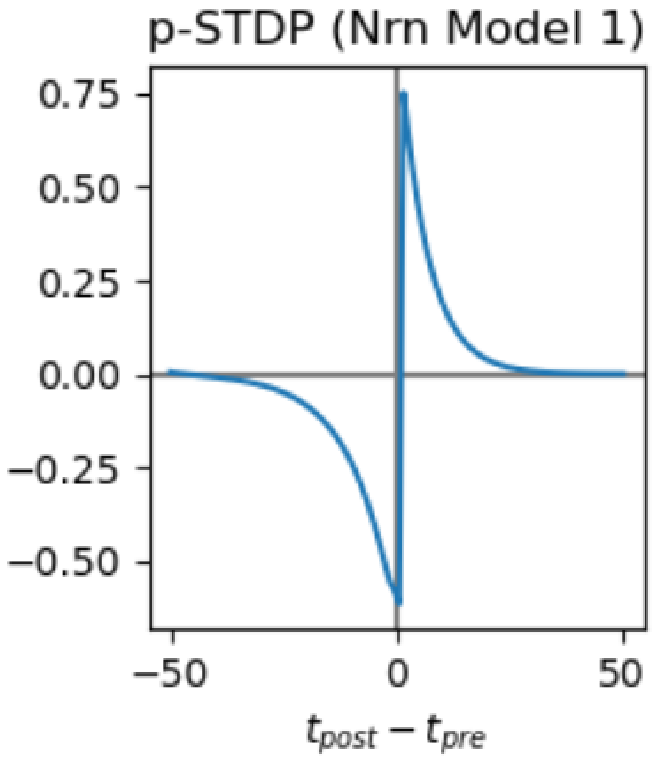

pSTDP1 rule can learn sequences

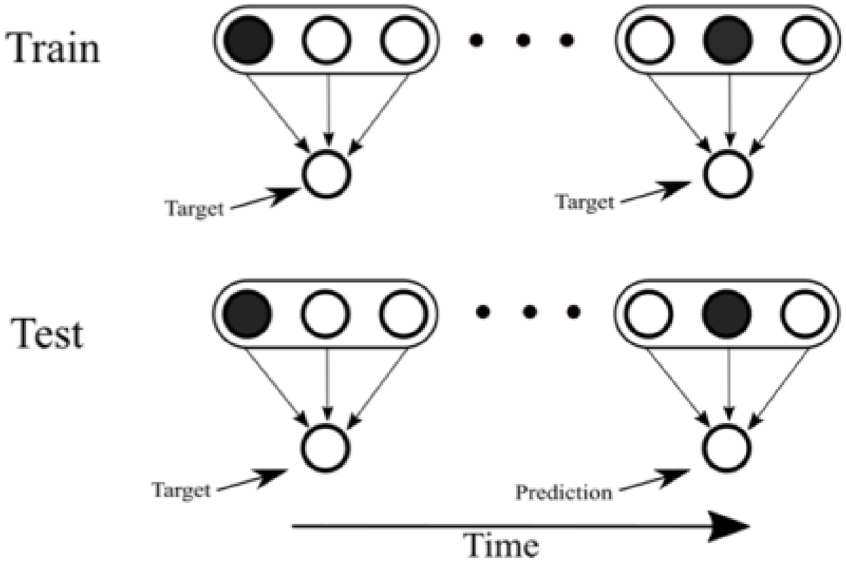

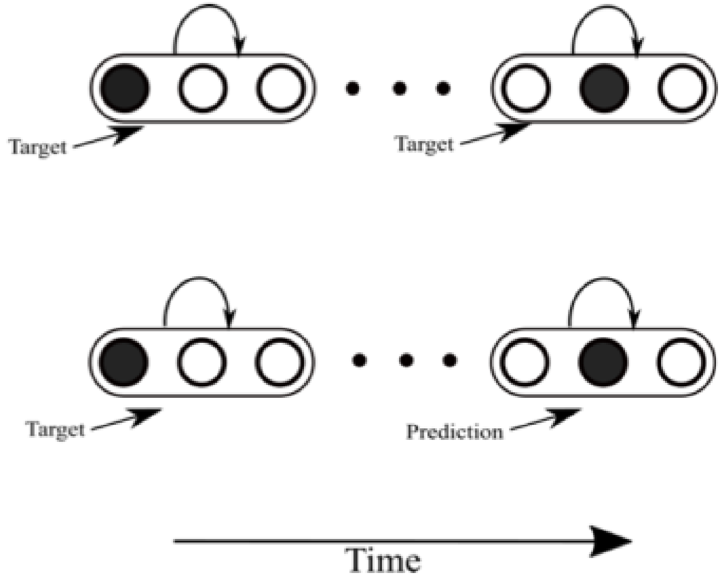

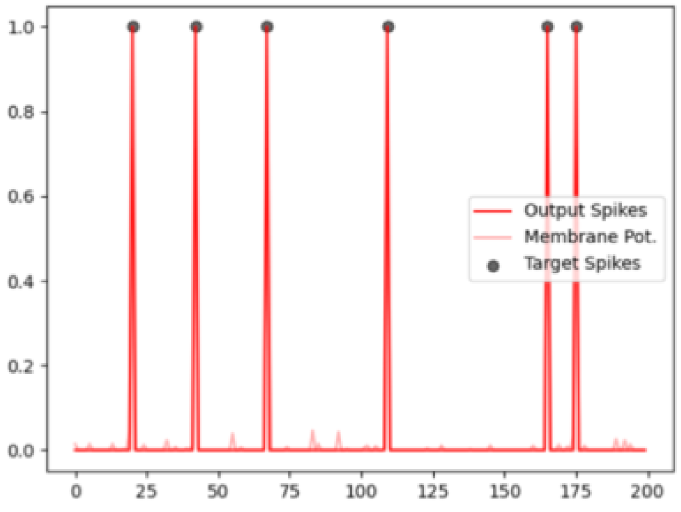

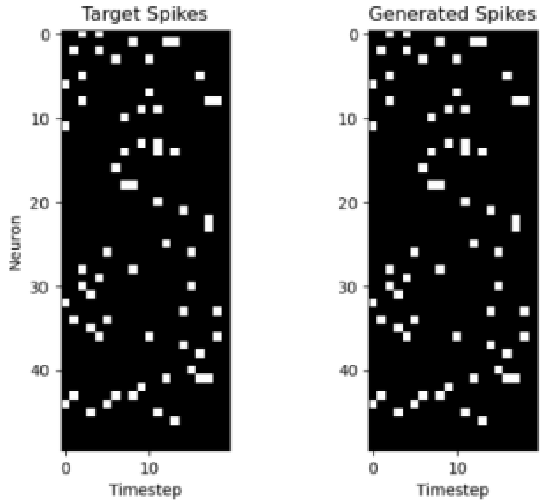

Derivation of pSTDP from Neuron Model 2

t: timestep

z: neuron spikes

W: synaptic weight matrix

α: decay rate

h: membrane potential

v: double trace

Neuron Model

$$h^{t+1} = \alpha h^t + W a^t - (\alpha h^t + W a^t) z^t$$

Eligibility Vectors

$$a^{t+1} = z^{t+1} + \alpha a^t$$

$$v^{t+1} = a^{t+1} + (1 - z^t) \alpha v^t$$

pSTDP2

$$-\frac{\partial L}{\partial W_{ij}} = -\frac{\partial \frac{1}{2}\|\Theta(h_j^{t+1}) - h_j^{t+1}\|^2}{\partial h_j^{t+1}} \sum_{t' \le t} \frac{\partial h_j^{t+1}}{\partial h_j^{t}} \cdots \frac{\partial h_j^{t'+1}}{\partial h_j^{t'}} \frac{\partial h_j^{t'}}{\partial W_j}$$

$$= -\frac{\partial L}{\partial e}\left(\frac{\partial e}{\partial z_j^{t+1}} \frac{\partial z_j^{t+1}}{\partial h_j^{t+1}} + \frac{\partial e}{\partial h_j^{t+1}}\right) v_i^t$$

$$\approx -(z_j^{t+1} - h_j^{t+1})(1 - a_j^{t+1} - 1) v_i^t$$

$\ldots$ (simplify)

$$= \underbrace{z_j^{t+1} v_i^t - W_{ij} a_i^t a_j^{t+1}}_{STDP-like} - \underbrace{b a_j^{t+1} v_i^t}_{modulation}$$

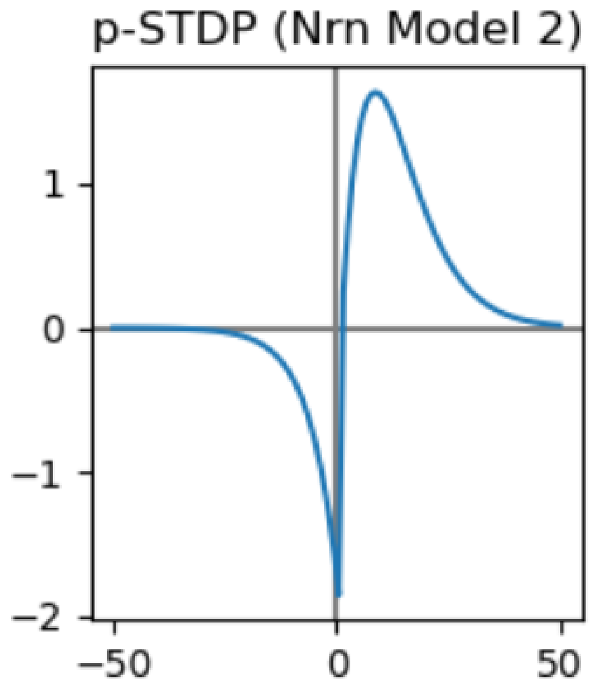

pSTDP2 can learn target time in delayed prediction task

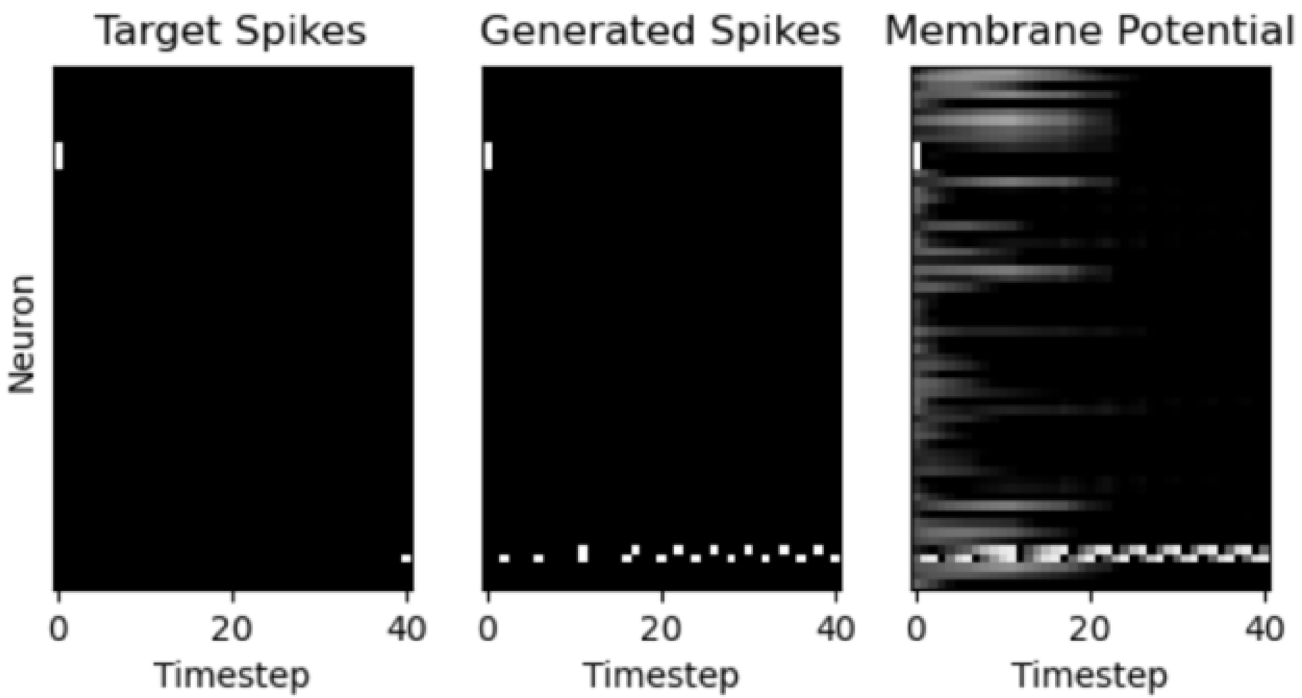

Figure 3: Spike raster for one input-output pair in the delayed spike prediction task. Here we see the network learned to produce the target spike at the desired time by propagating immediately generated spikes in the target and neighboring neuron, which combined signals to produce the target spike at the desired time.

Summary

Learning rules derived via eligibility propagation and a predictive coding loss function

- Resemble standard STDP rules
- Can learning spiking sequences with delays

Pre synaptic neurons are trying to predict the behavior of postsynaptic neurons at future timesteps

We suggest that STDP not only optimizes synapses for prediction, but does so by preforming credit assignment though time using something akin to eligibility propagation.