Los Alamos
NATIONAL LABORATORY
EST. 1943

Implementing Backpropagation for Learning on Neuromorphic Spiking Hardware

Alpha Renner, Forrest Sheldon, Anatoly Zlotnik, Louis Tao,
Andrew Sornborger

ETH zürich

LABORATORY DIRECTED RESEARCH & DEVELOPMENT

NICE 2021
March 17, 2021

LA-UR-21-22463

ASC

NNSA National Nuclear Security Administration

NNSA National Nuclear Security Administration

Preliminary Results: MNIST

MNIST 10x10 | Dense 300 | Dense 10

- 3 layer network (300 hidden units)
- Binary and downsampled by 2

- Training – 96%, Validation – 92%

- 14 Loihi timesteps per training sample
- Inference after 3 timesteps

- 1000 FPS, 1 ms/sample
- 0.3 mJ/sample
- Energy-delay product = 0.3µJs

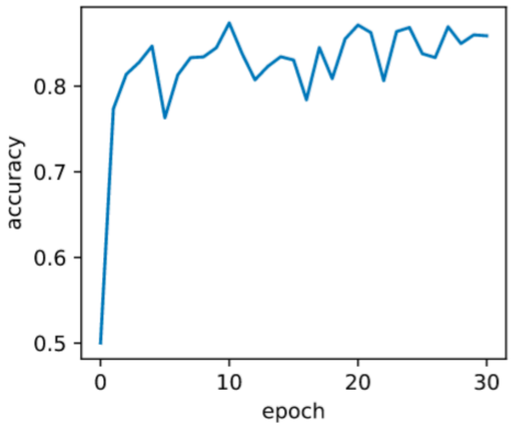

Accuracy over epochs. Validation set.

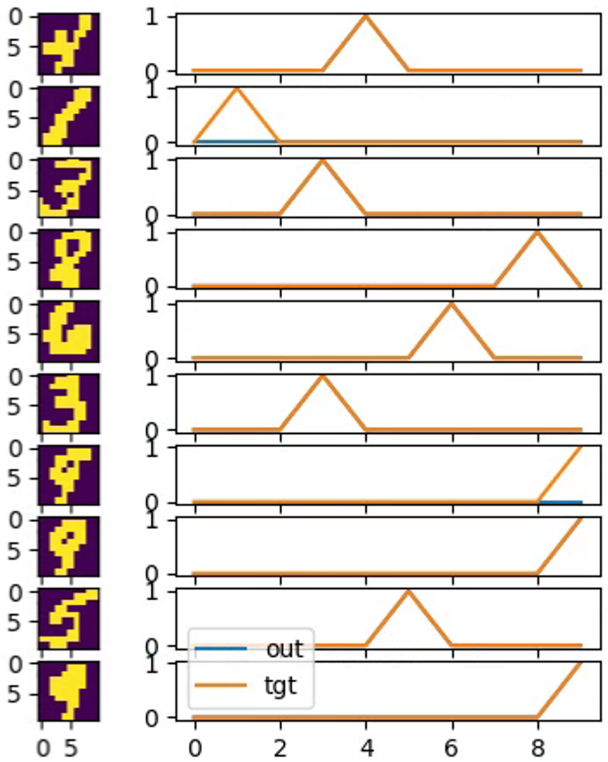

Spiking Backprop

Neural and network mechanisms for implementing backprop:

- *Synfire-gated synfire chain(s)*
- *Short-term memories*
- *Push-me pull-you pairs for encoding real numbers*
- *Gating of ReLU thresholded activity*
- *Gating of derivative of ReLU (theta function) activity via SGSC*
- *Implementation of Hadamard product via pulse-gating*
- *Simultaneous gating of graded information to pre- and post-synaptic neuronal populations for hebbian synaptic update (turning learning on and off via pulse-gated control)*

Sornborger, Andrew, Louis Tao, Jordan Snyder, and Anatoly Zlotnik. "A Pulse-gated, Neural Implementation of the Backpropagation Algorithm." In *Proceedings of the 7th Annual Neuro-inspired Computational Elements Workshop*, pp. 1-9. 2019.

Alpha Renner, Forrest Sheldon, Louis Tao, Anatoly Zlotnik, Andrew Sornborger. "A Pulse-gated, Spiking Neural Implementation of the Backpropagation Algorithm." to appear In *Proceedings of the 78h Annual Neuro-inspired Computational Elements Workshop*, pp. 1-9. 2020.

Backpropagation Algorithm

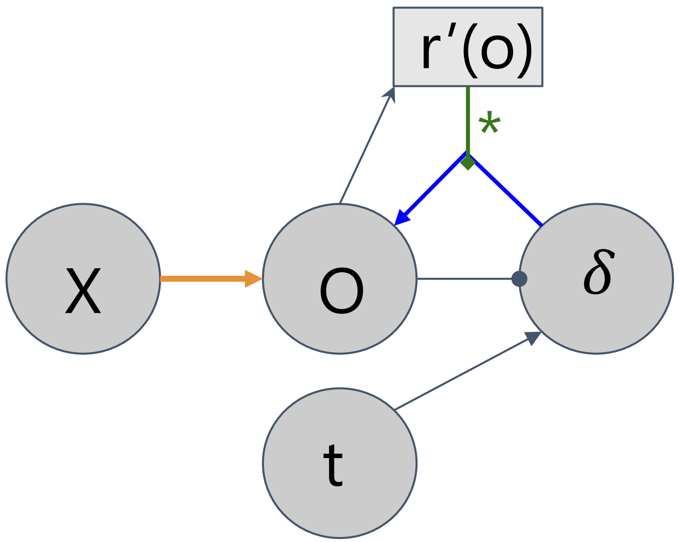

Update for a single neuron:

$$\Delta w_0^{ij} = \delta_a^i \cdot x^j \cdot r'(z^i)$$

"error" Input Derivative of activation function

The learning mechanism in detail

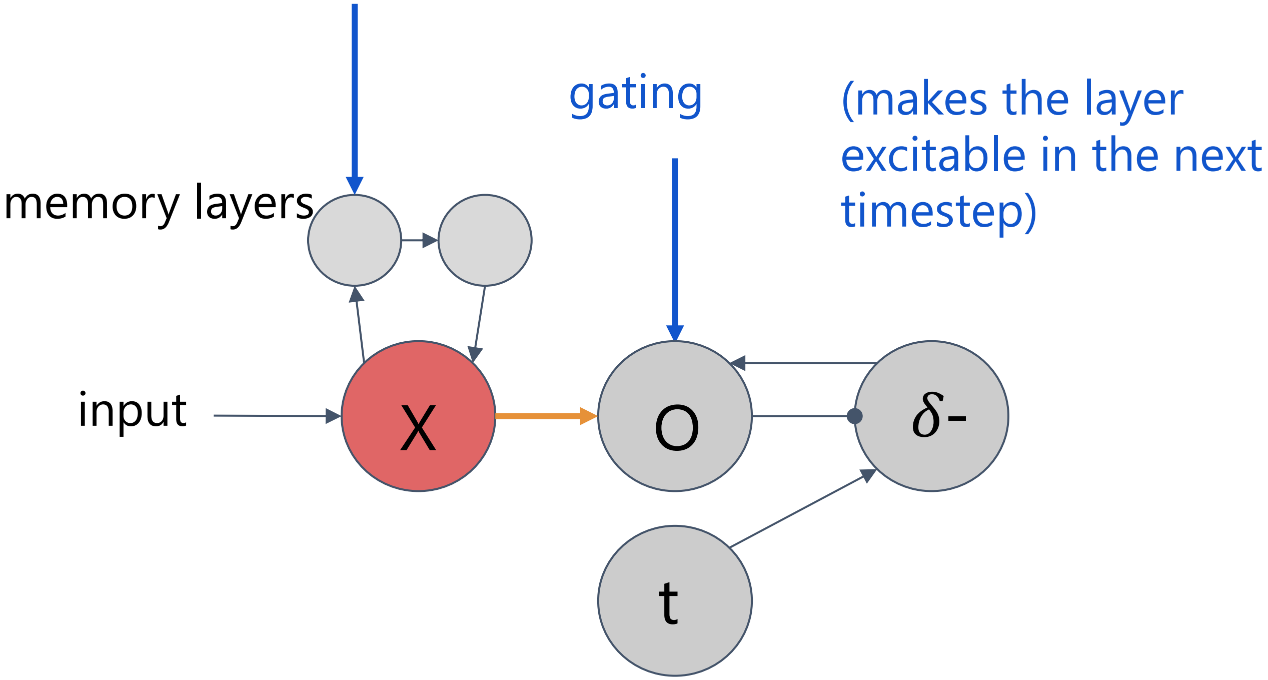

The learning mechanism in detail

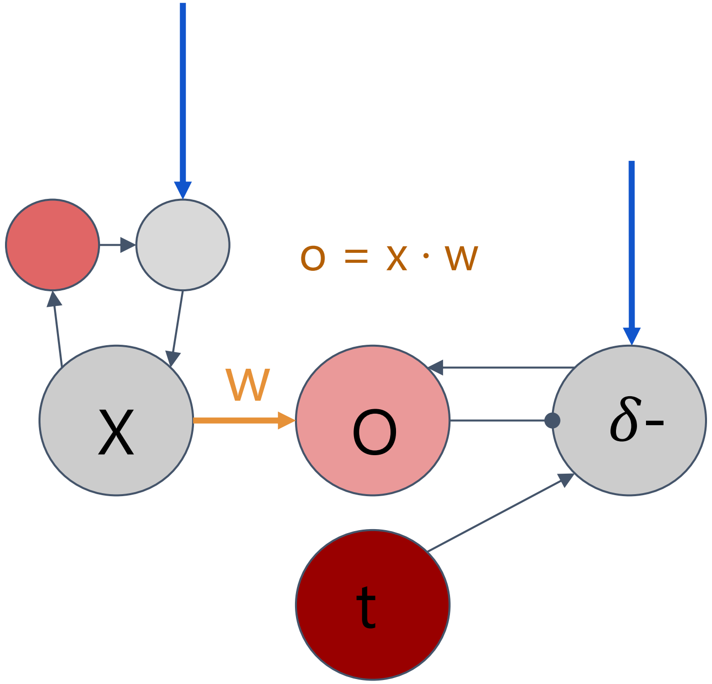

The learning mechanism in detail

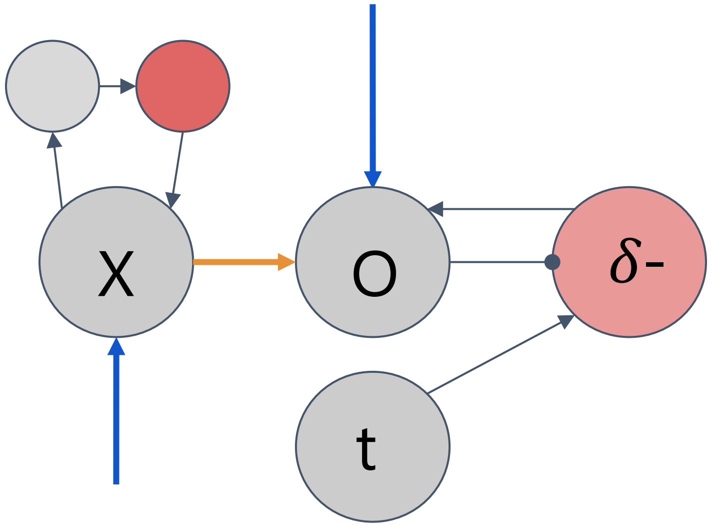

$\delta^{-} = t - o$

The learning mechanism in detail

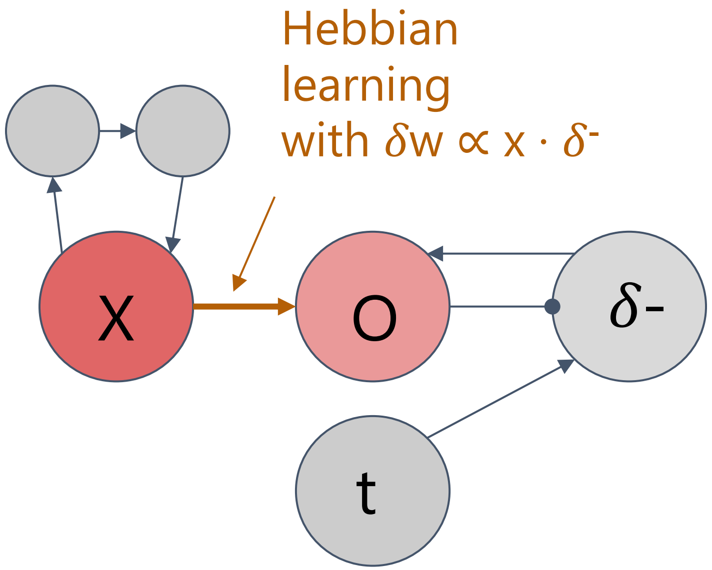

Note:

This is just a simplified visualization, the actual δw is:

$\delta w \propto \delta \cdot x \cdot r'(o)$

The learning mechanism in detail

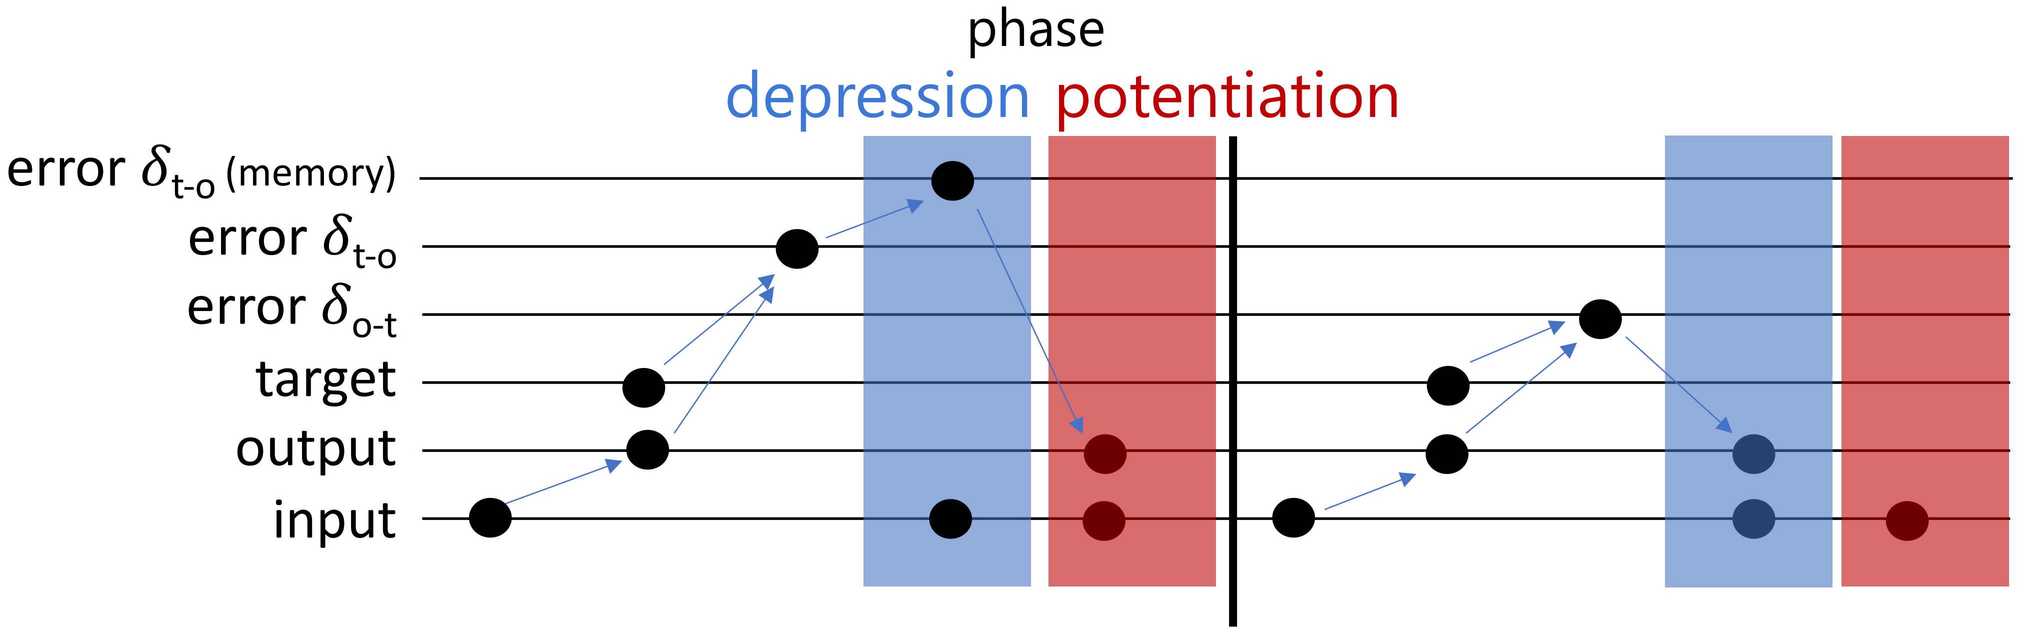

Target > Output
→ Hebbian learning in potentiation phase

Target < Output
→ Hebbian learning in depression phase

Multilayer Backpropagation Algorithm

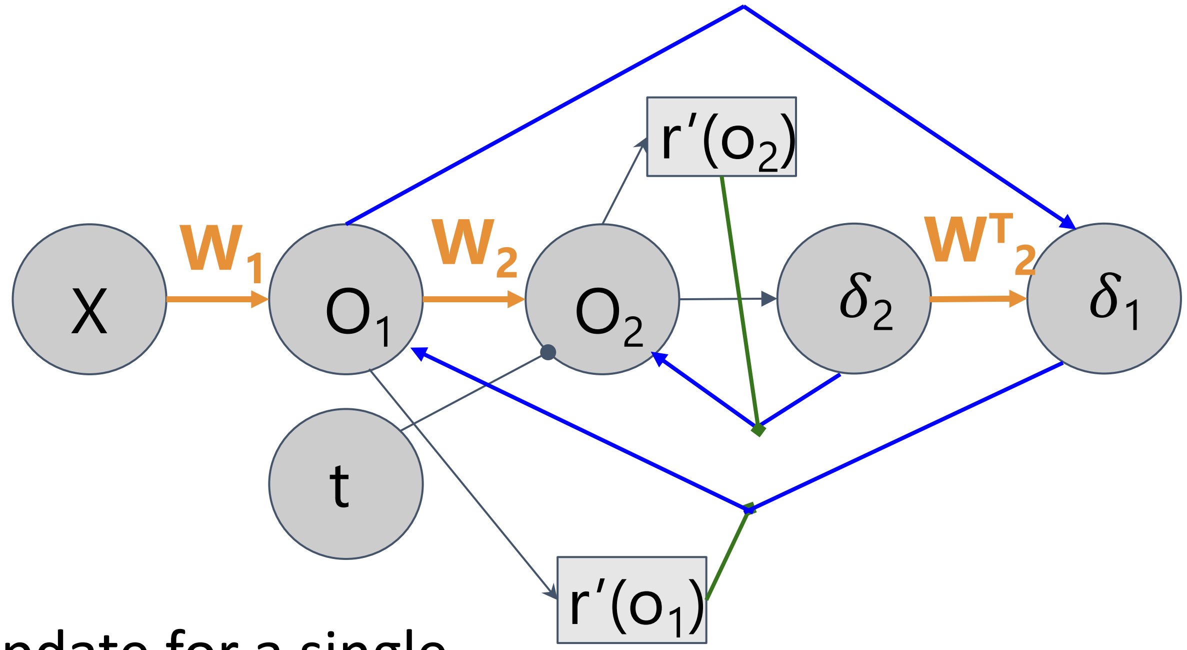

Feedback alignment or transpose of weight matrix is needed

Update for a single neuron:

$$\Delta w_0^{ij} = \delta_a^i \cdot x^j \cdot r'(z^i)$$

$$\delta_a^i = w_1^{(:,j)T} \cdot \delta_z$$

δ of next layer

Spiking Backprop

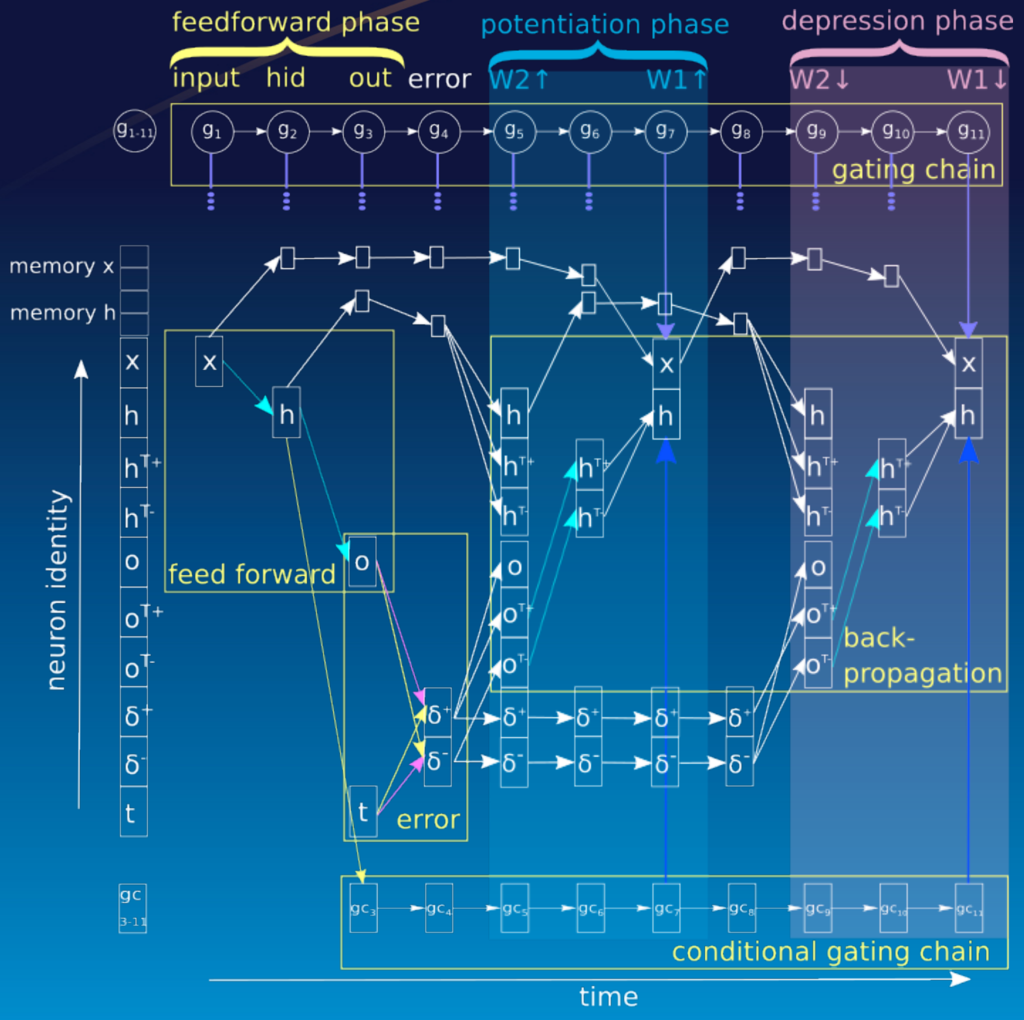

- all:all plastic (all:all on population level)
- 1:1 excitatory (1:1 on population level)
- 1:1 excitatory (all:all on population level)
- 1:1 inhibitory (all:all on population level)
- 1:all excitatory gating (1:all on population level)
- 1:1 excitatory gating (1:all on population level)

global third factor positive

global third factor negative

Alpha Renner, Forrest Sheldon, Louis Tao, Anatoly Zlotnik, Andrew Sornborger. "A Pulse-gated, Spiking Neural Implementation of the Backpropagation Algorithm." In *Proceedings of the 78th Annual Neuro-inspired Computational Elements Workshop*, pp. 1-9. 2020.

Spiking Backprop - XOR

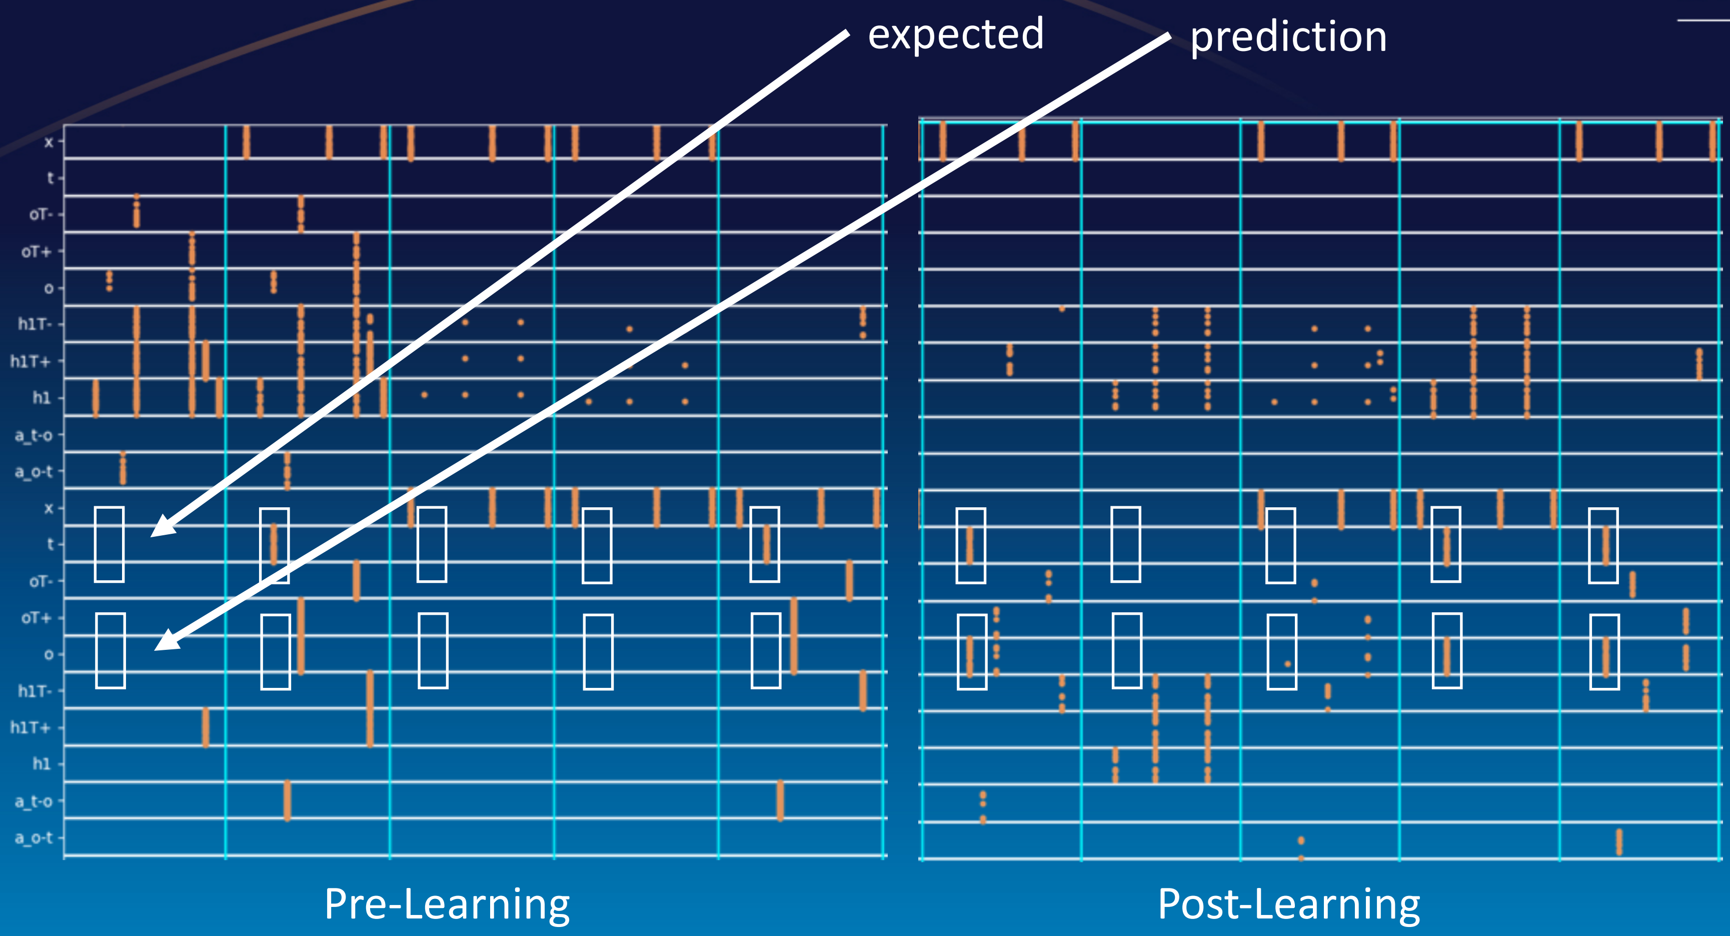

Alpha Renner, Forrest Sheldon, Louis Tao, Anatoly Zlotnik, Andrew Sornborger. "A Pulse-gated, Spiking Neural Implementation of the Backpropagation Algorithm." to appear In *Proceedings of the 78th Annual Neuro-inspired Computational Elements Workshop*, pp. 1-9. 2020.

Preliminary Results: MNIST

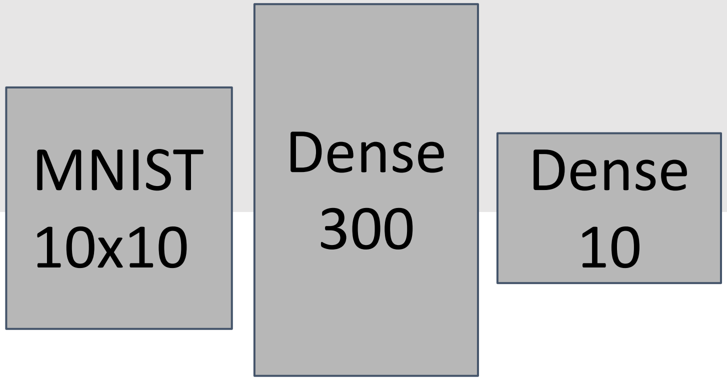

- Training – 96%, Validation – 92%

- 14 Loihi timesteps per training sample
- Inference after 3 timesteps

- 1000 FPS, 1 ms/sample
- 0.3 mJ/sample
- Energy-delay product = 0.3μJs

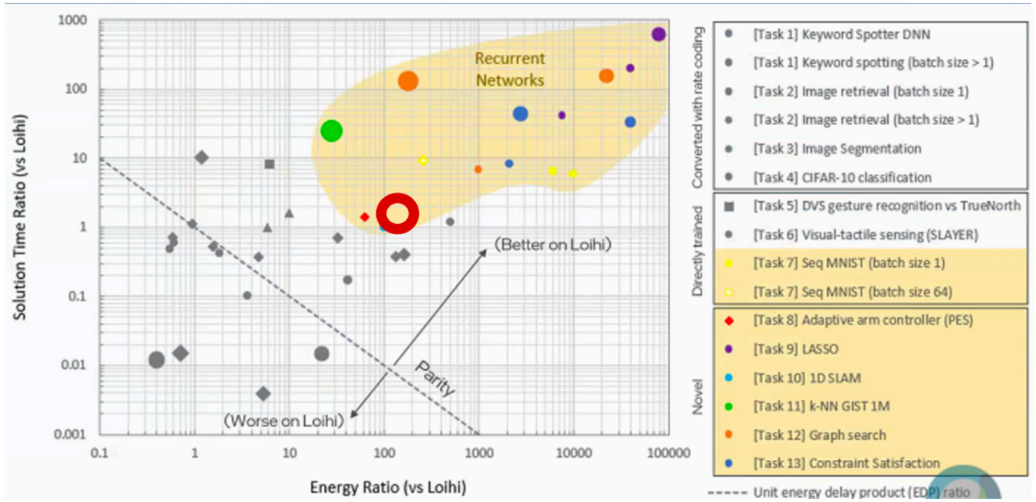

- 3 layer network (300 hidden units)
- Binary and downsampled by 2

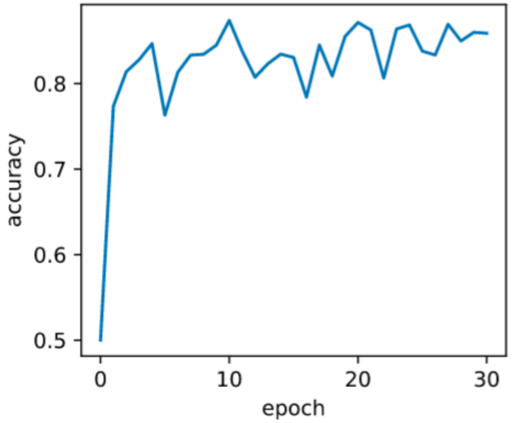

Accuracy over epochs. Validation set.

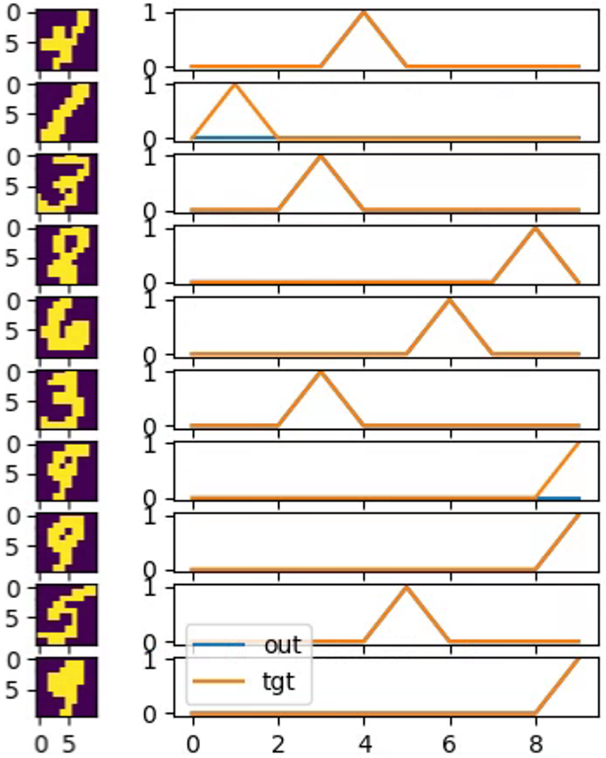

Conclusion

- Proof of principle of pulse-gated, spiking Backpropagation algorithm
- We use
  - Synfire-gated synfire chains to precisely route and remember values
  - Hebbian learning in 2 phases (depression and potentiation) to update weights (solution for the 3-factor dilemma of other approaches)
  - Binary coding (saves time, neurons and spikes) or population rate coding to encode values and enable a ReLU-like activation function
- MNIST: 1000 FPS, Energy-delay product – 0.3µJs

Questions?