Training Neuromorphic Systems for Scientific Applications

Gangotree Chakma

University of Tennessee

Katie Schuman

Oak Ridge National Laboratory

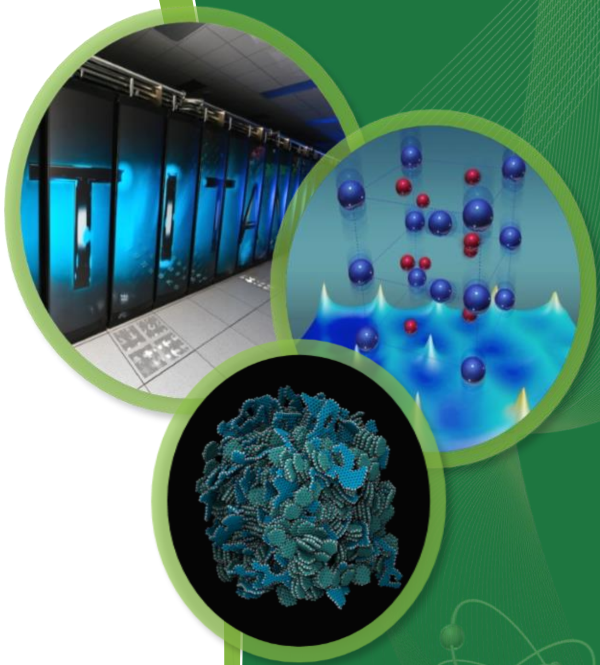

TENN LAB
NEUROMORPHIC
ARCHITECTURES. LEARNING. APPLICATIONS.

Spiking Recurrent Neural Networks

- There are key elements for a spiking recurrent neural network that must be defined:
  - What is the topology of the network?
    - How many neurons?
    - Level of connectivity?
    - Recurrent connections?
  - What should the parameters of the network be?
    - Synaptic weights
    - Neuron thresholds
    - Delays

Possible Training Algorithms

	Topology Defined or Trained?	Delays Utilized?	Delays Defined?	Weights Defined?	Training Time	Demonstrated Broad Applicability
Back-propagation	Some	No	Some	Yes	Medium	Some
STDP	Some	Yes	No	Yes	Medium	No
Evolutionary Approaches	Yes	Yes	Yes	Yes	Slow	Yes
Liquid State Machines	Some	Yes	Random	Random	Medium	Some

Good	Okay	Potentially Bad

Spiking Recurrent Neural Networks for Neuromorphic

- Neuromorphic systems often support spiking networks, along with variable delays and recurrent connections.
- To fully utilize such a system, programmable weights, delays, and topology should be utilized.
- Smaller networks and those with less activity may correspond to lower energy or power usage.
- For any given application, it is not always clear how to adapt the algorithm.
  - Some algorithms are relatively inflexible for non-classification problems.

Neuroscience-Inspired Dynamic Architectures (NIDA)

- Spiking neural network embedded in 3D space.
- Simple neuron (integrate-and-fire) and synapse implementation.
- Flexible structure.

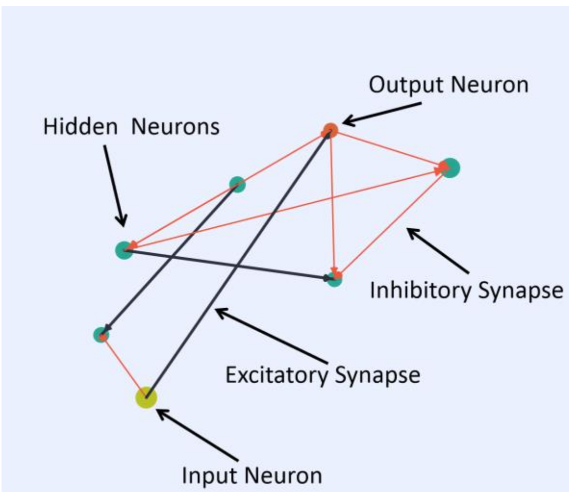

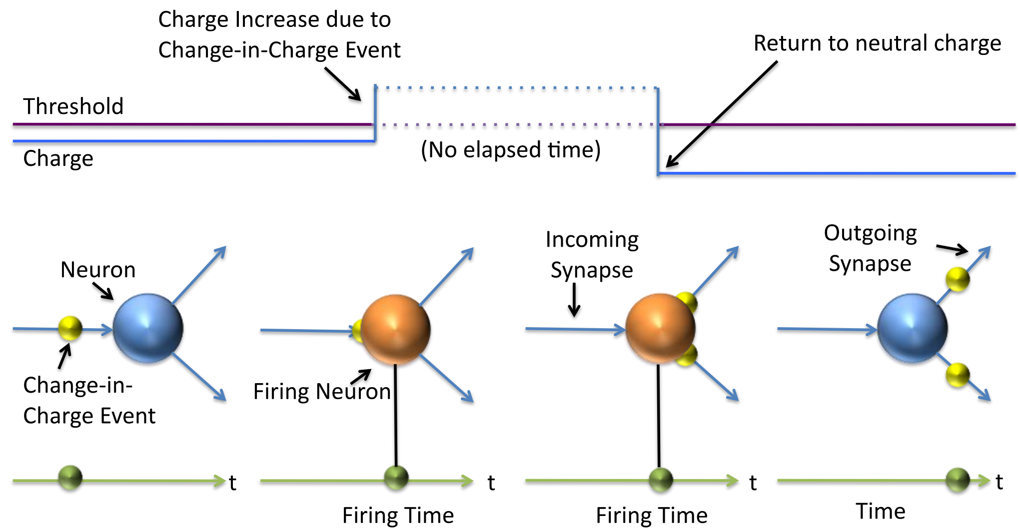

TENNLAB NEUROMORPHIC ARCHITECTURES. LEARNING. APPLICATIONS.

OAK RIDGE National Laboratory

Memristive DANNA (mrDANNA)

- Mixed analog/digital implementation.
  - Mixed signal analog neurons.
  - Each synaptic weight is implemented with two memristors.
- Lower energy, better scaling than digital implementations.
- Fabricating with 65nm cmos 10lpe node in collaboration with CNSE, SUNY PI, Albany, NY.

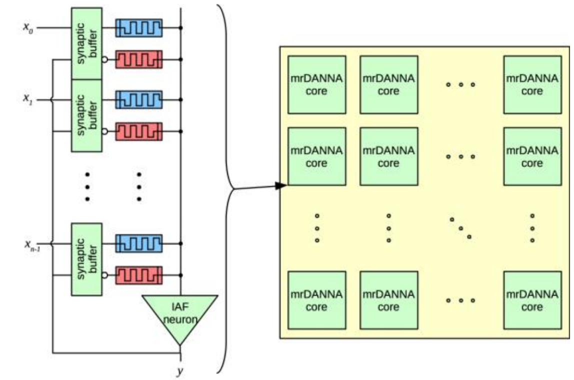

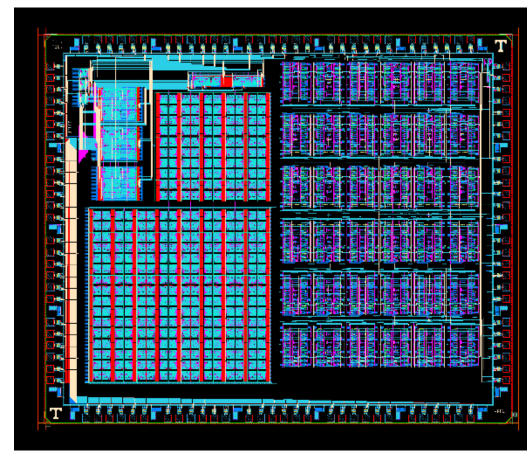

Evolutionary Optimization for Neuromorphic Systems (EONS)

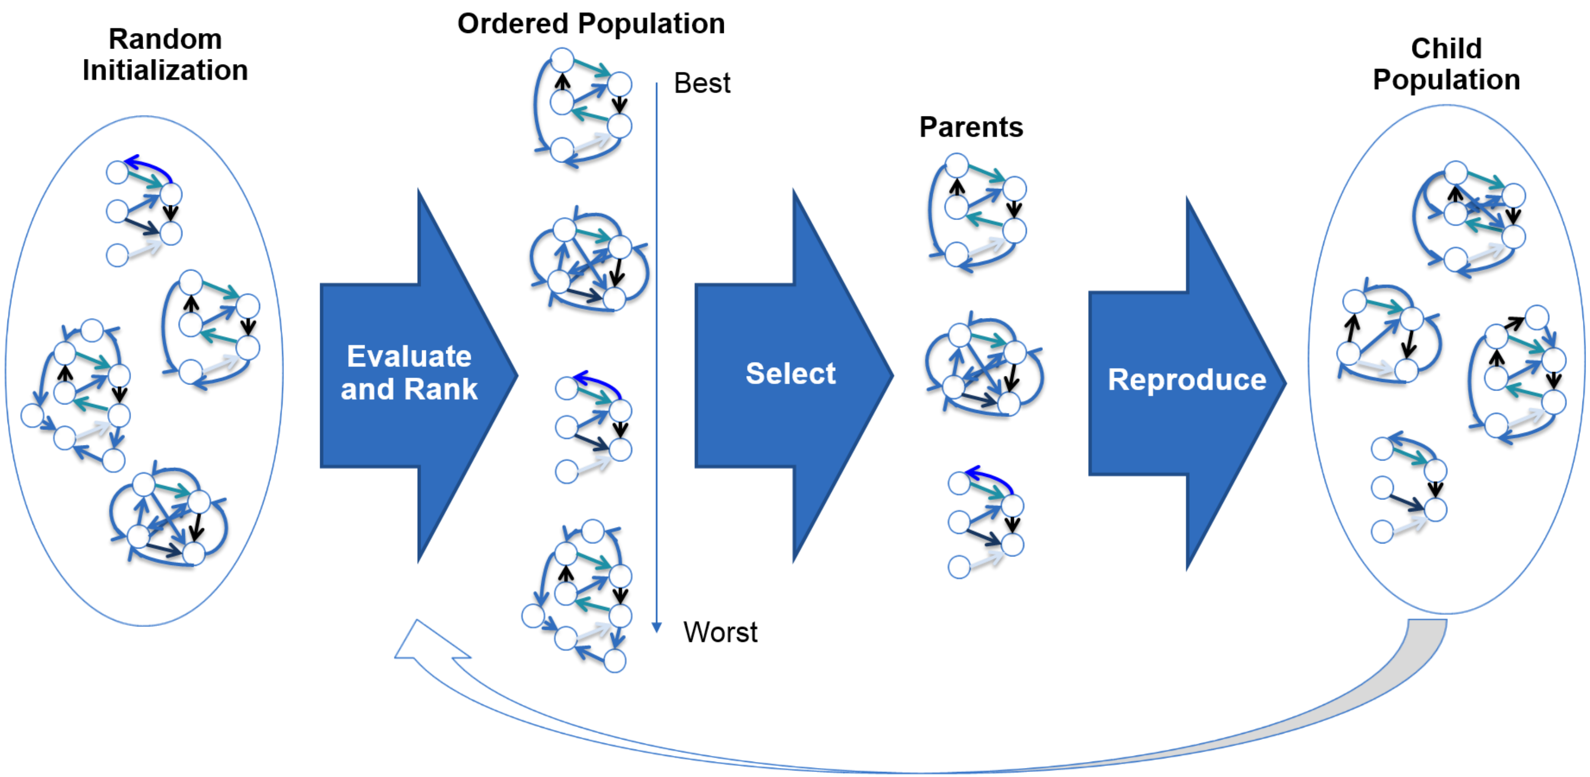

TENN LAB NEUROMORPHIC ARCHITECTURES. LEARNING. APPLICATIONS.

OAK RIDGE National Laboratory

Reservoir Computing Approach

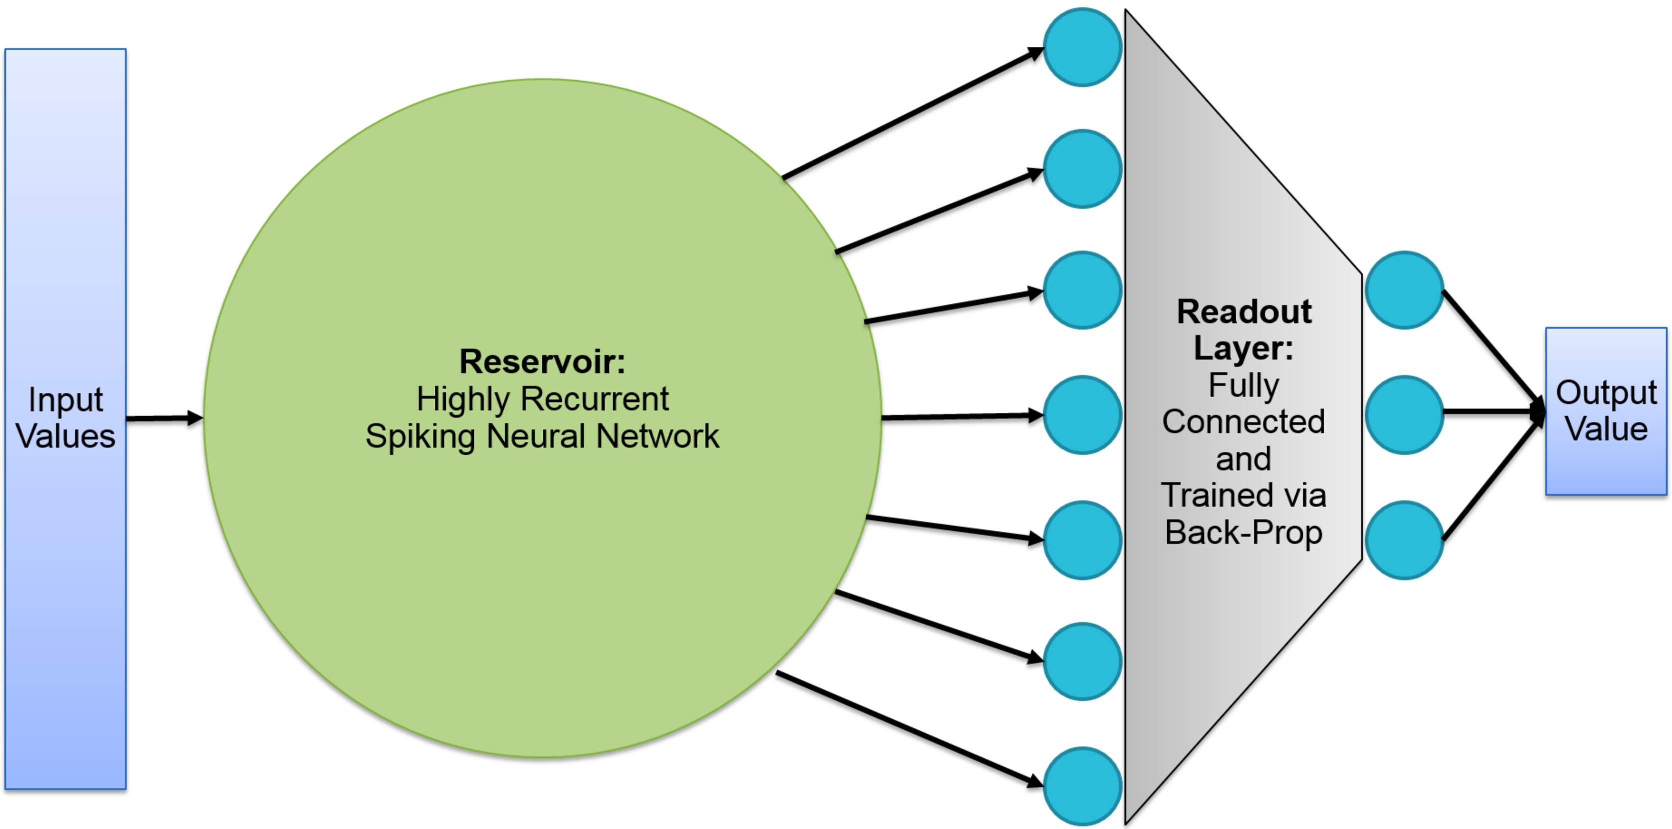

Data from MINERvA (Main Injector Experiment for v-A)

- Neutrino scattering experiment at Fermi National Accelerator Laboratory
- The detector is exposed to the NuMI (Neutrinos at the Main Injector) neutrino beam.
- Millions of simulated neutrino-nucleus scattering events were created.
- Classification task is to classify the horizontal region where the interaction originated.

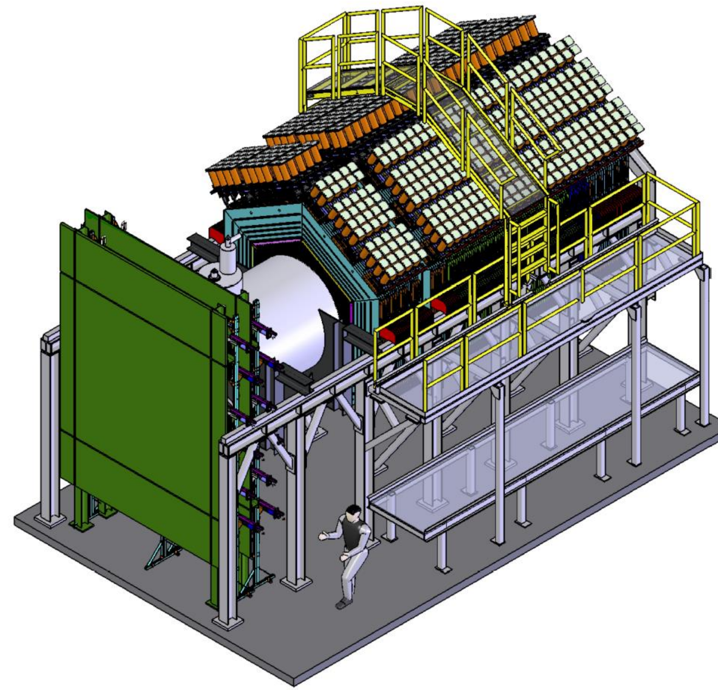

MINERvA Detector

Source: A. Terwilliger, et al. Vertex Reconstruction of Neutrino Interactions using Deep Learning. IJCNN 2017.

Two Data Inputs Types (Three Views)

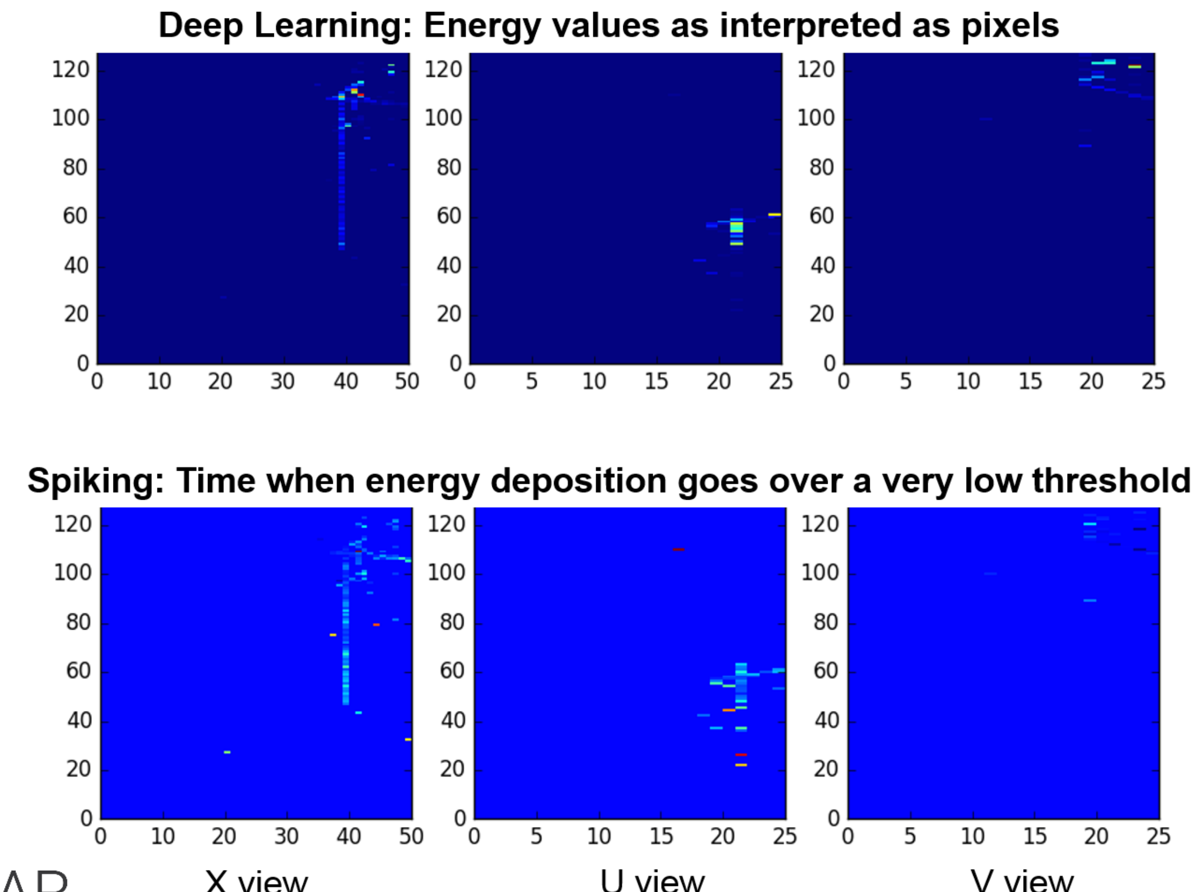

Best Results: Single View

x-view (127x50)	conv1 (8x3)	pool1 (2x1)	conv2 (7x3)	pool2 (2x1)	conv3 (6x3)	pool3 (2x1)	conv4 (6x3)	pool4 (2x1)	fc1 (196)	drop out	fc2 (98)	drop out	fc3 (11)	classification

Convolutional Neural Network Result: ~80.42%

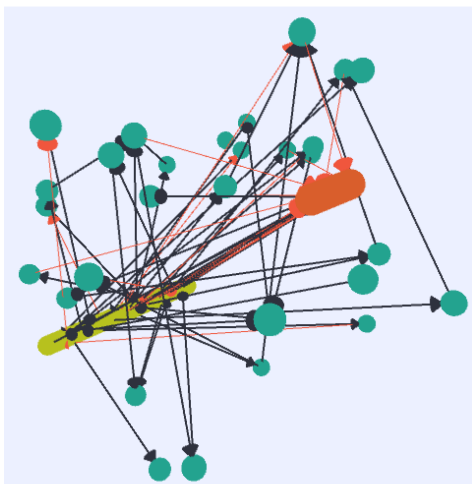

- 90 neurons, 86 synapses
- Estimated energy for a single classification for mrDANNA implementation: 1.66 μJ

Spiking Neural Network Result: ~80.63%

Source for CNN results: A. Terwilliger, et al. Vertex Reconstruction of Neutrino Interactions using Deep Learning. IJCNN 2017.

Preliminary Results

Number of Nodes	Maximum Classification Accuracy after One Hour
100	70.82%
1000	72.6%
10000	79.11%

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

TENN LAB
NEUROMORPHIC
ARCHITECTURES. LEARNING. APPLICATIONS.

OAK RIDGE
National Laboratory

Best Network Performance

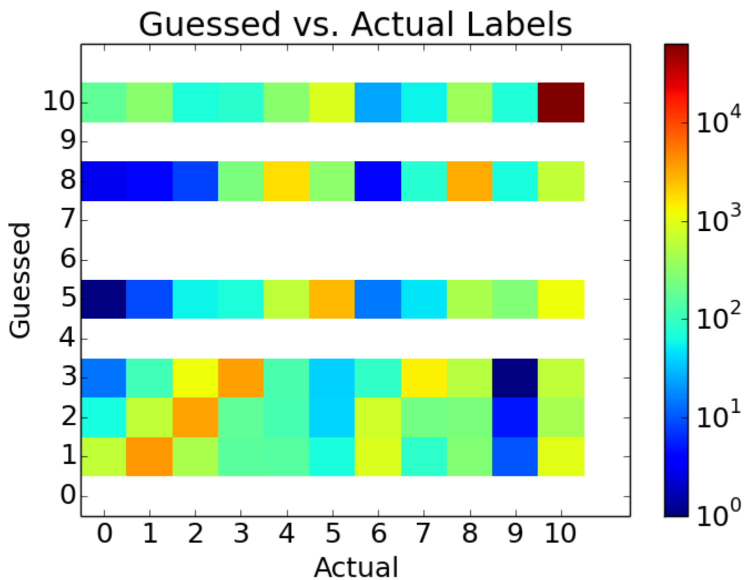

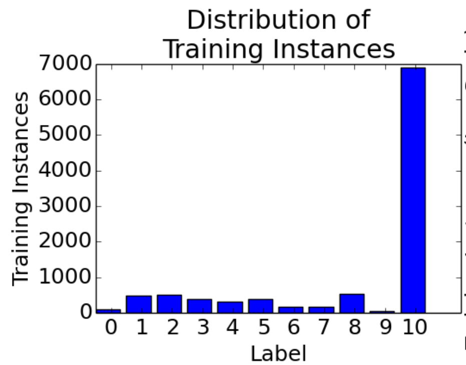

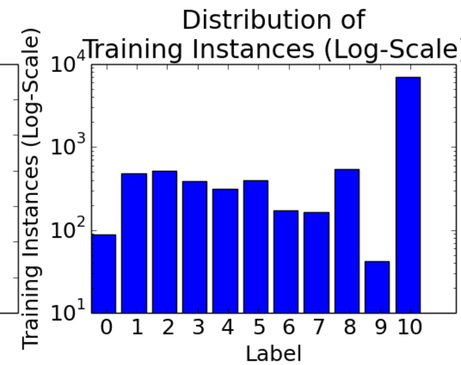

Updated Fitness Function Results

Punish networks that don't choose every label at least once.

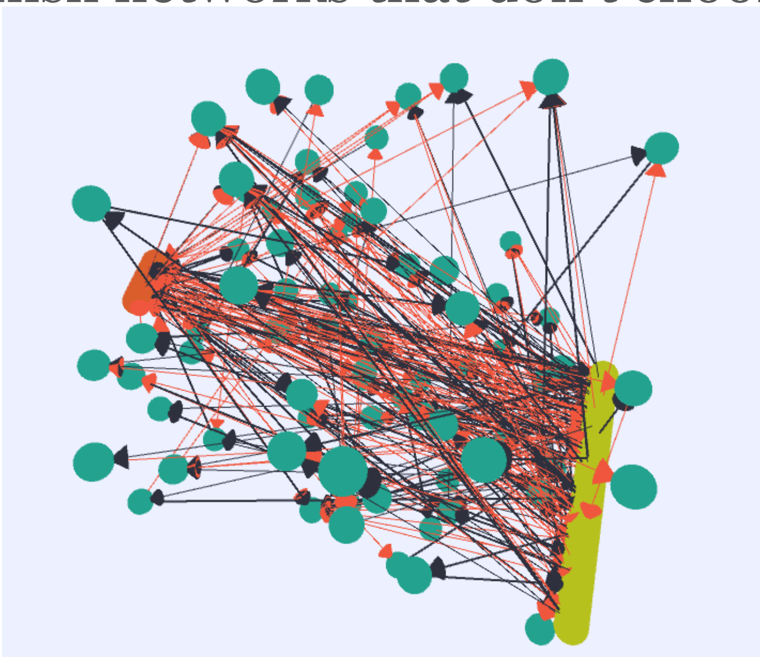

- 140 neurons
- 355 synapses
- Training accuracy: **86.57%**
- Test accuracy: **83.62%**

TENN LAB NEUROMORPHIC ARCHITECTURES. LEARNING. APPLICATIONS.

OAK RIDGE National Laboratory

Preliminary Reservoir Computing Results

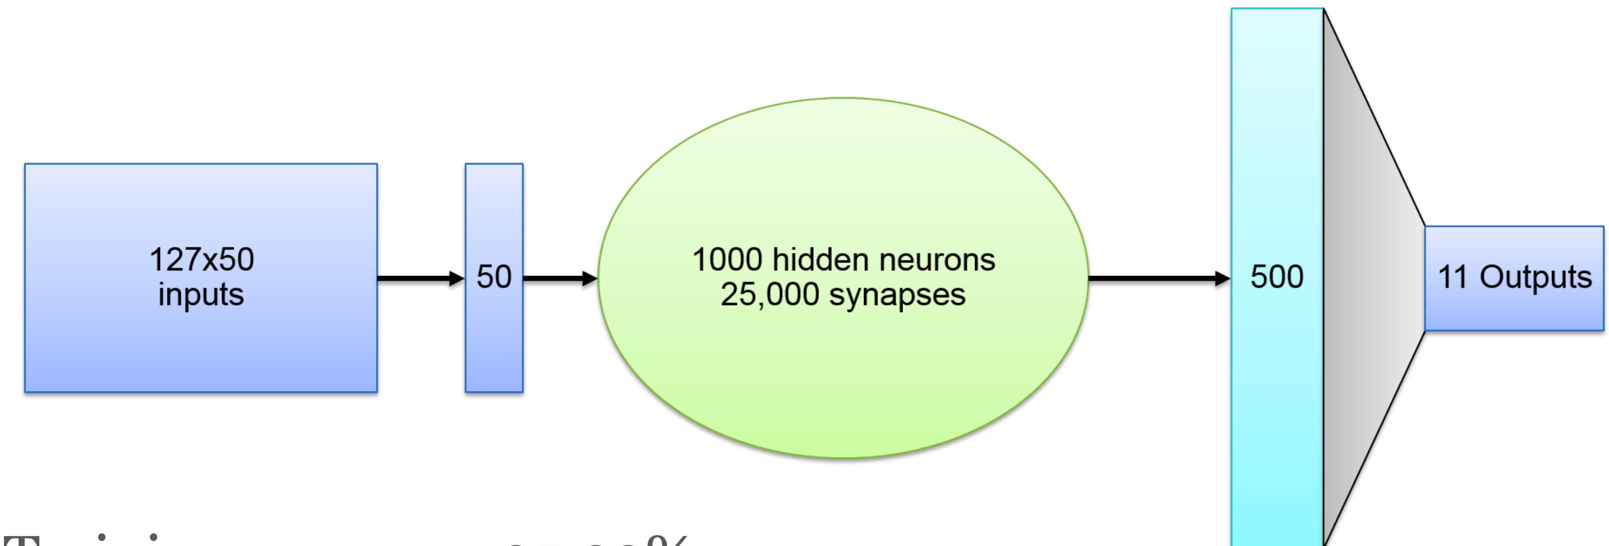

- Training accuracy: 95.33%
- Test accuracy: 81.03%

Comparison of Different Methods

	Testing Accuracy	Traditional Network Size (Number of weights)	Spiking Network Size	Approximate Training Time
Convolutional Neural Network trained using Back-Propagation	80.42%	~45,000 weights	0	~2 days on a GPU
Spiking Neural Network trained using EONS	83.62%	0	140 neurons, 355 synapses	~4 days on a 50 node cluster
Liquid State Machine with spiking liquid	81.03%	5500 weights	1550 neurons, 25000 synapses	~2 days on a single machine (Almost all in simulation of the network)

Conclusions and Discussion

- We achieved the best performance with the spiking neural network trained using evolutionary optimization, but liquid state machines also show promising results.
- All three approaches produced comparable accuracies.
- There is a tradeoff in the resulting network size and the training time to achieve such a network.
- Back-propagation and liquid state machines may be able to be trained faster with less resources, but the resulting networks tend to be much larger than those designed by EONS.
- If the resulting network is to be deployed into hardware, smaller networks will have a smaller footprint and will also likely required less power.

Acknowledgements

Contact: Katie Schuman (schumancd@ornl.gov)