A Study of Complex Deep Learning Networks on High Performance, Neuromorphic, and Quantum Computers

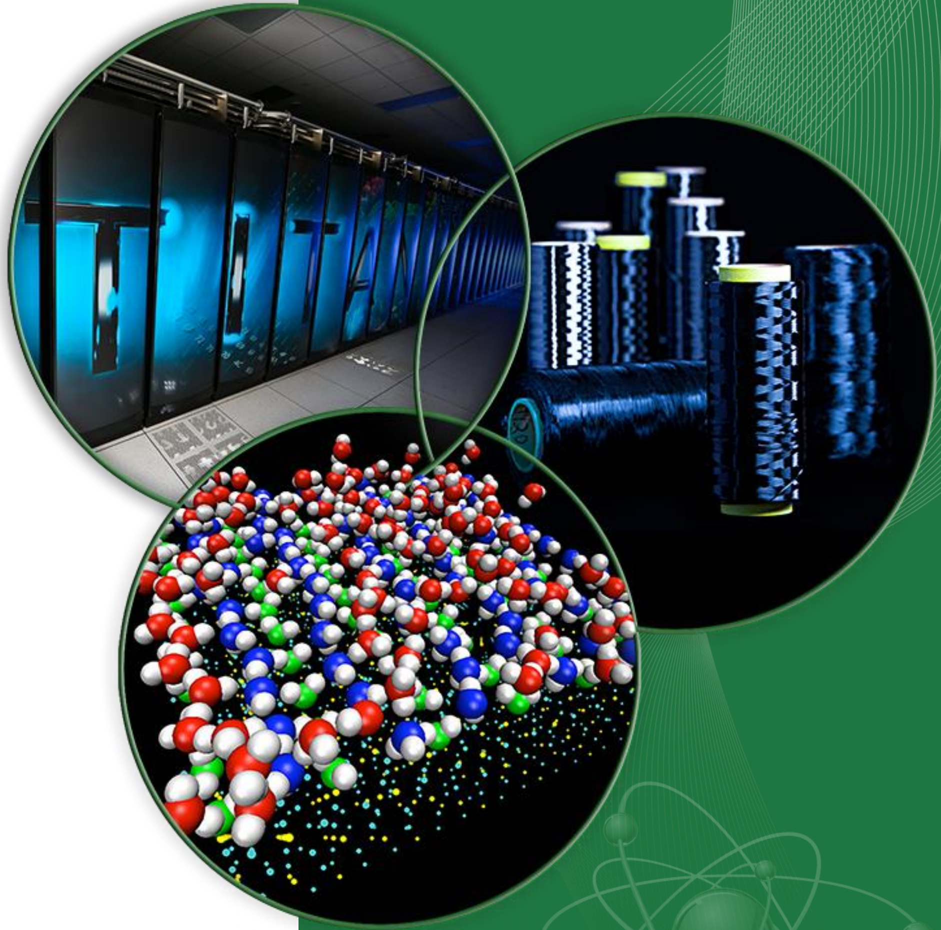

Thomas E. Potok, Ph.D.

Computational Data Analytics Group

Oak Ridge National Laboratory

ORNL is managed by UT-Battelle
for the US Department of Energy

U.S. DEPARTMENT OF ENERGY | Office of Science

OAK RIDGE National Laboratory

Motivation and Goals

- Scientific data is increasingly large and complex, making new discovery difficult
- Can deep learning and novel architectures provide a way of aiding scientific discovery?

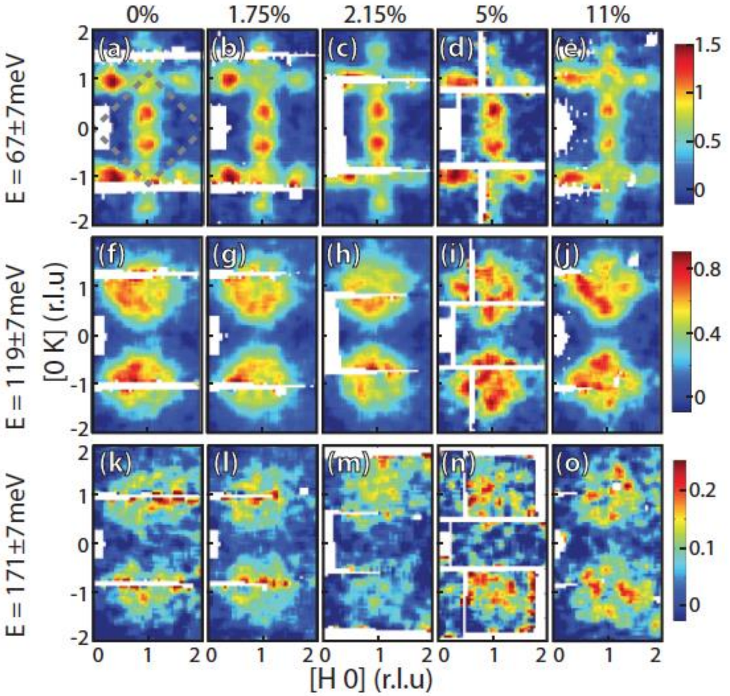

Each pixel represents a 512x512 matrix of values from Spallation Neutron Source at ORNL

Deep Learning Performance

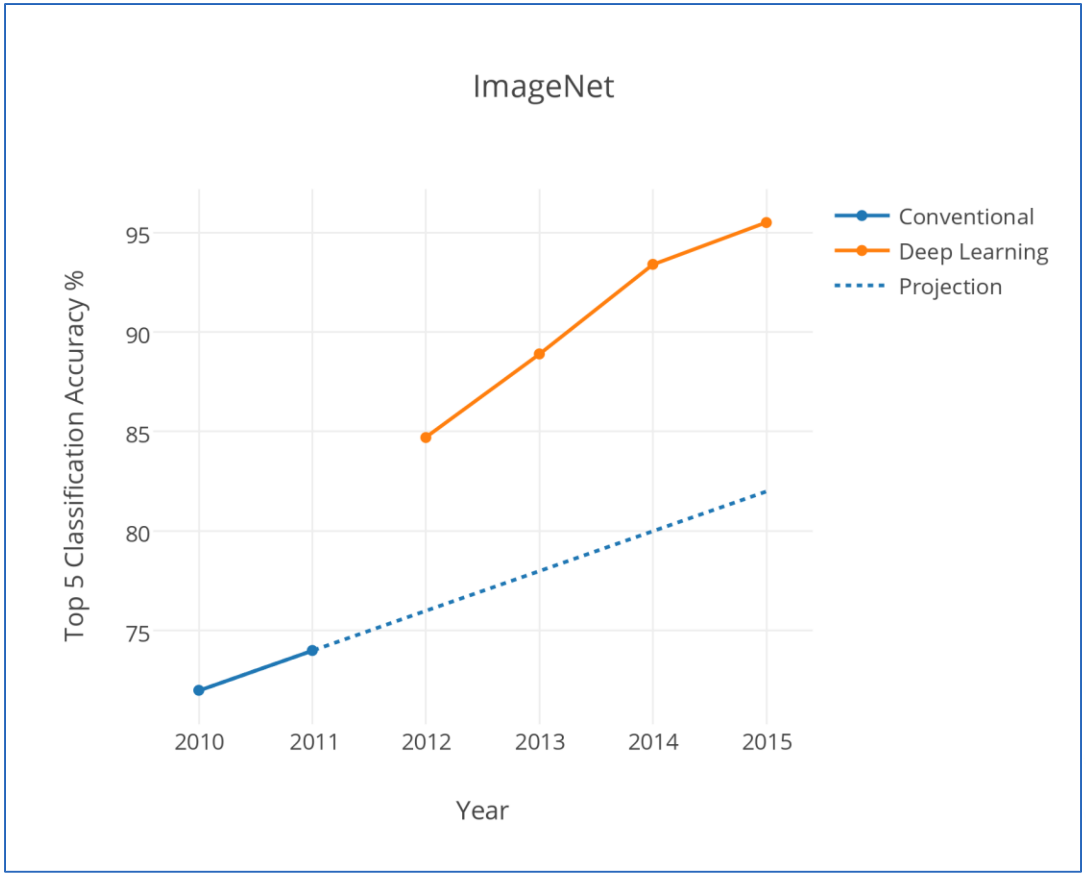

- What could you do with
  - HPC,
  - Quantum
  - Neuromorphic?

OAK RIDGE National Laboratory

Motivations for novel architectures

- Deep Learning network topologies are layered due to computability
  - Does a complex topology offer better results?
- Deep Learning Hyper-parameters are hard to tune
  - Can a CNN be quickly tuned for a new scientific datasets?
- Deep Learning Trained models are hard to deploy
  - Can models be deployed on or near scientific instruments?

Hypothesis

Scientific Data

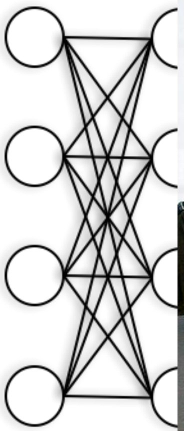

Comple
Topology

ardware
plementation

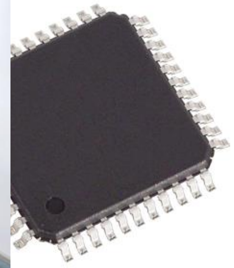

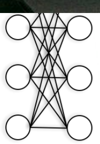

Auto Tuned Hyper
Parameters

Scientific Instrument

Methods

Quantum

HPC

Neuromorphic

Scientific Data

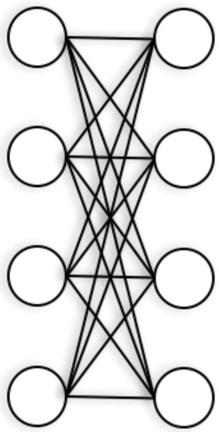

Complex Topology

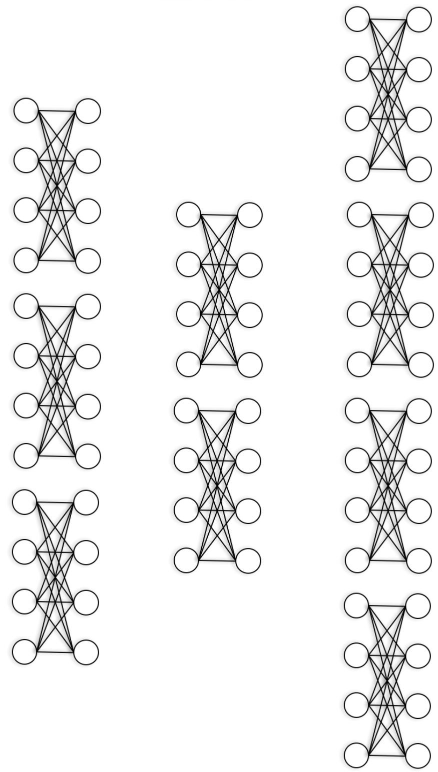

Auto Tuned Hyper Parameters

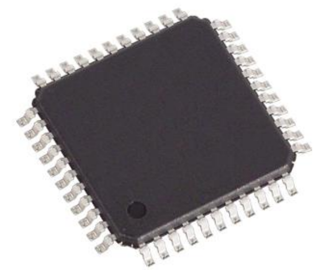

Hardware Implementation

Scientific Instrument

OAK RIDGE National Laboratory

Rational

Quantum	• Hints at speed up of certain hard problems
HPC	• Computational parallelization
Neuromorphic	• Native neural networks, very low power

Experimental Goals

Quantum

- Test feasibility of complex topologies (intra-layer connections)
- USC Lockheed Martin Quantum Computation Center D-Wave

HPC

- Test evolutionary algorithms to auto-tune hyper-parameters
- ORNL's Titan

Neuromorphic

- Test the ability to represent neural network models in low power hardware
- UTK Memristive Nida System

How do we compare the architectures?

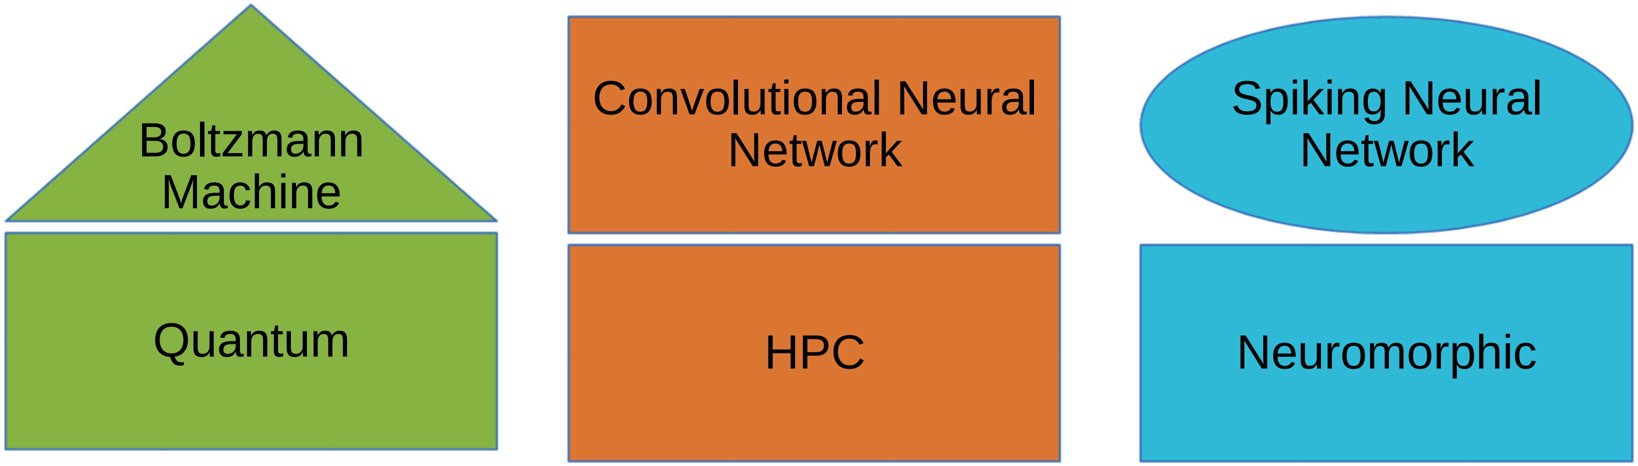

- The MNIST database contains 60,000 training images and 10,000 testing images
- Very small images size, 28x28
- http://yann.lecun.com/exdb/mnist/

Oak Ridge National Laboratory

Common Ground

- Quantum Physics
- Computer Science
- Electrical Engineering
- Neuroscience

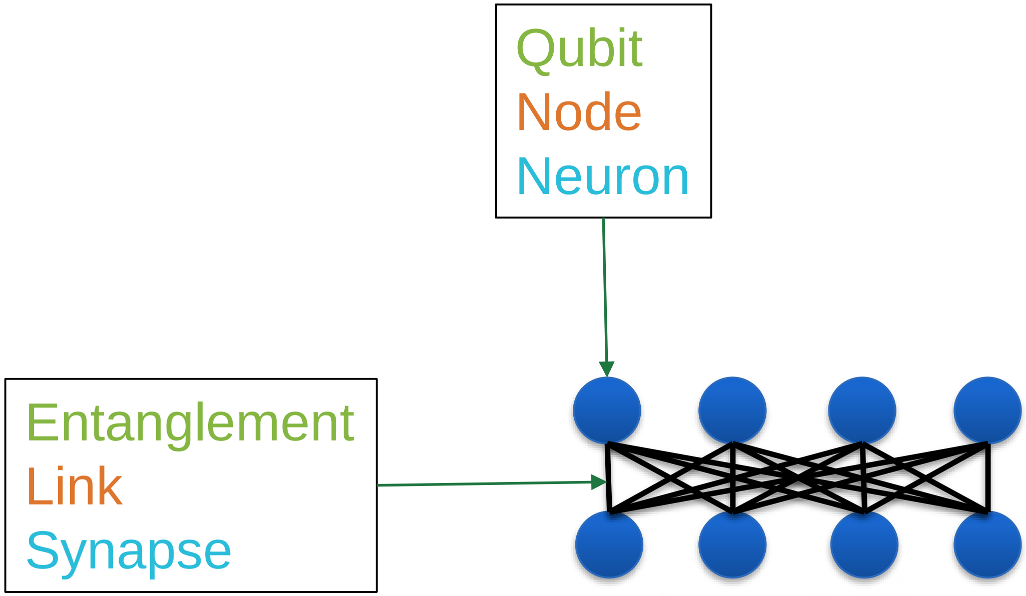

Chimera Network

Limited Boltzmann Machine Network

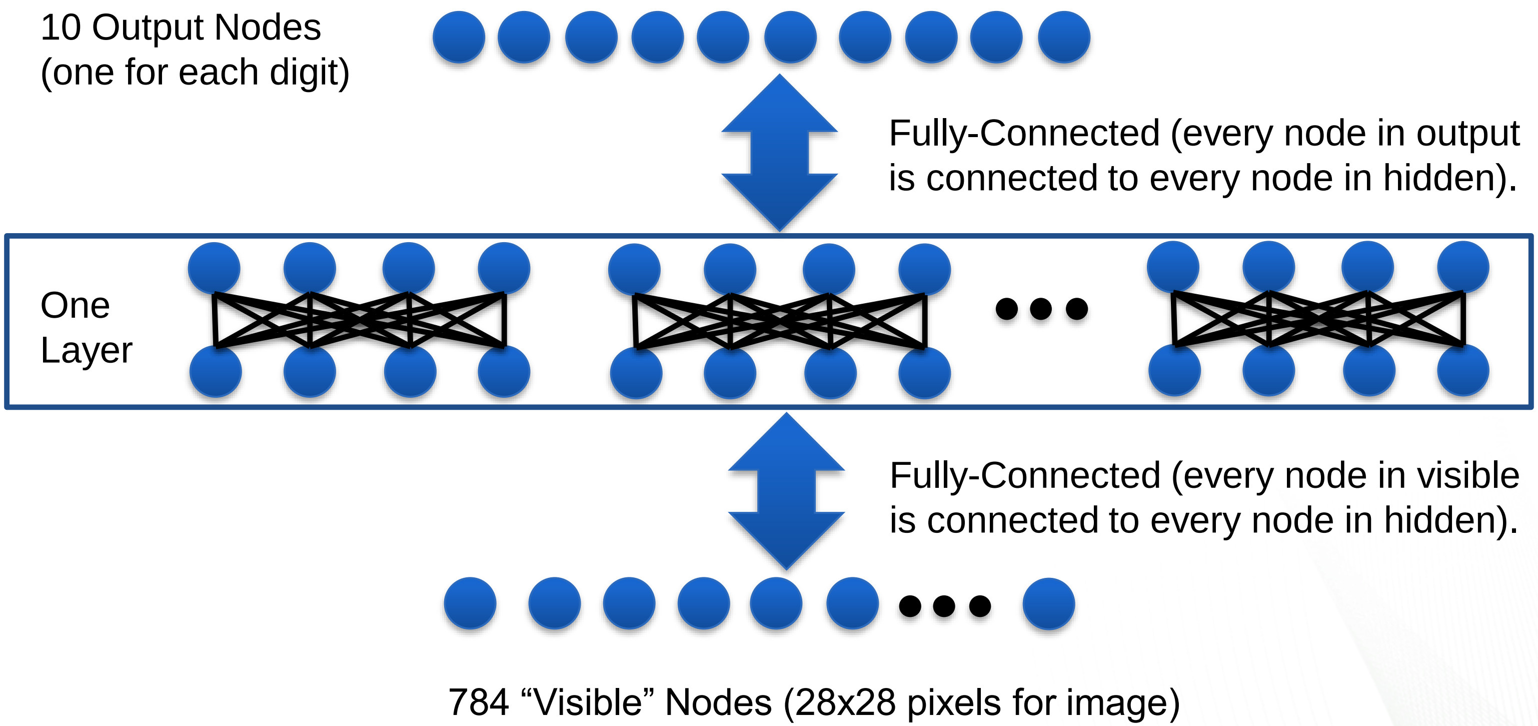

Quantum Results – USC/ISI D-Wave

Complex topology learns and provides better results than no intra-layer connection

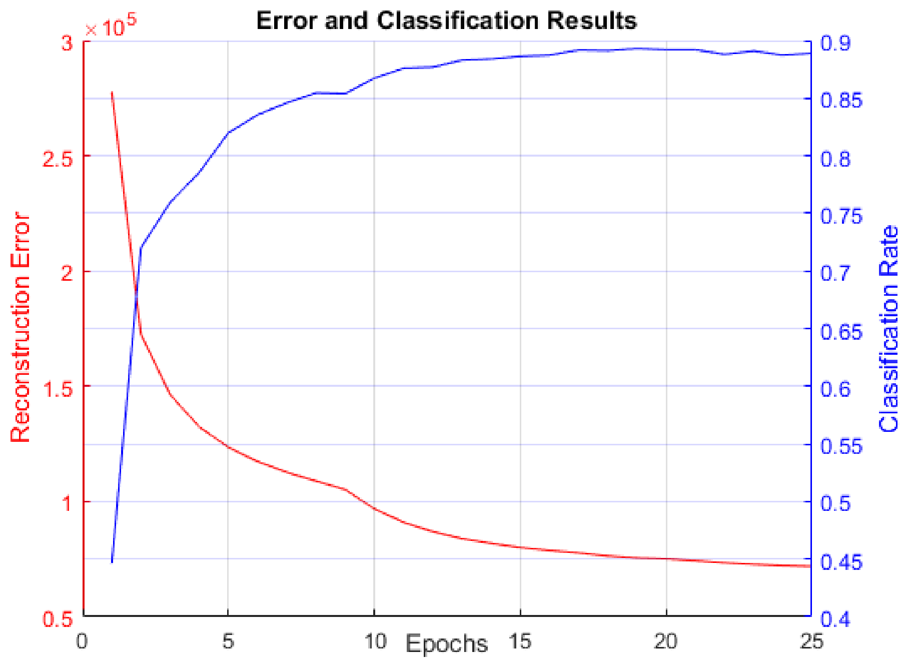

Limited Boltzmann Machine learning from training examples

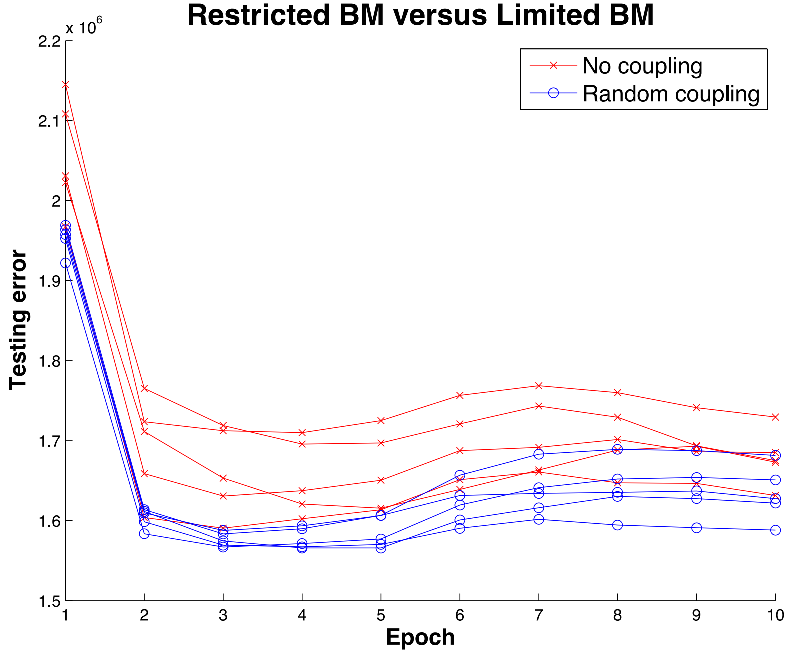

Limited Boltzmann Machine more accurate than restricted BM

OAK RIDGE National Laboratory

HPC Results – 500 nodes on Titan

Demonstrates an effective way of auto tuning a CNN

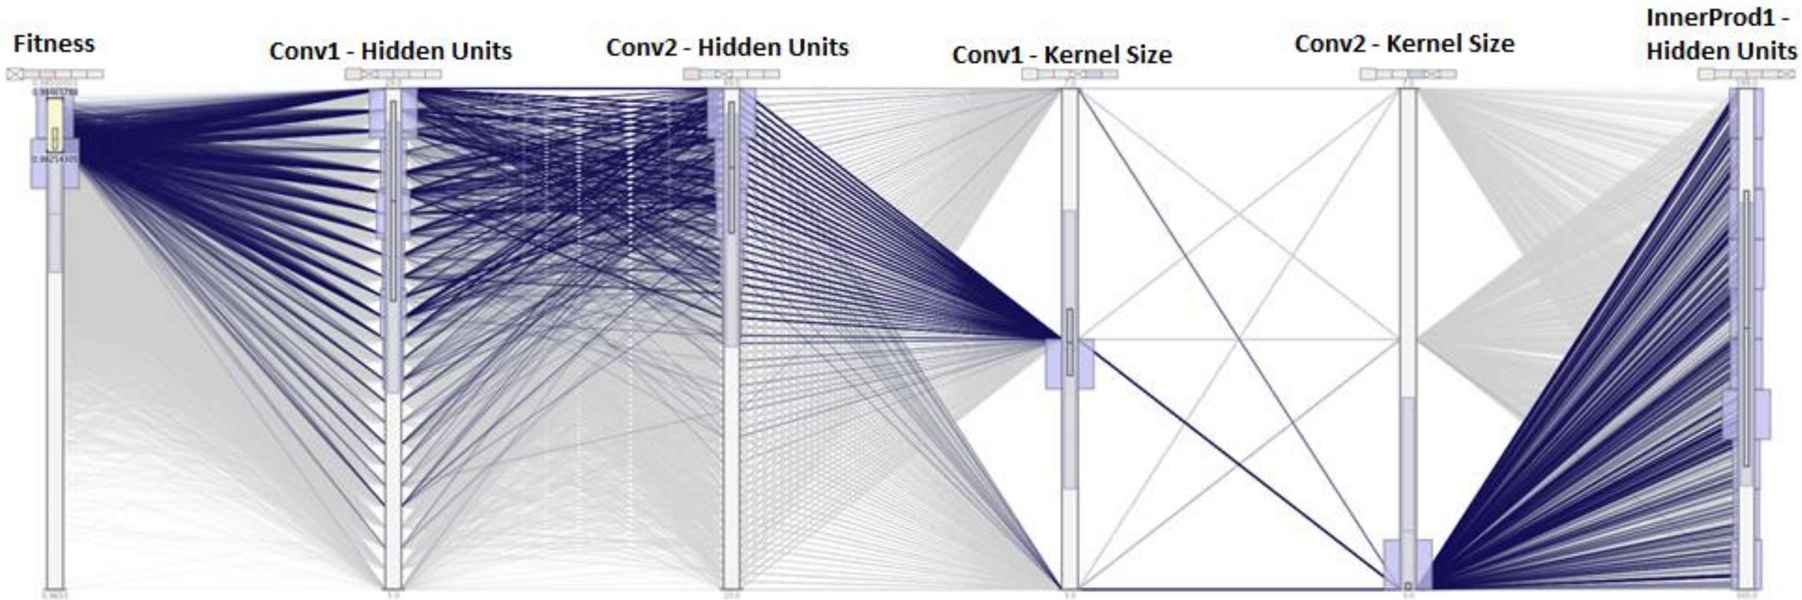

MNIST

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |

Inner Product Layer (10 outputs)

Inner Product Layer (500, 590 hidden units)

Pool2 Layer (Max, K=2)

Conv2 Layer (50, 67 hidden units, K=5, 3)

Pool1 Layer (Max, K=2)

Conv1 Layer (20, 29 hidden units, K=5, 5)

MNIST Training Data

Neuromorphic – ORNL/UT NIDA network on Memristor

20X more energy-efficient than their CMOS counterparts

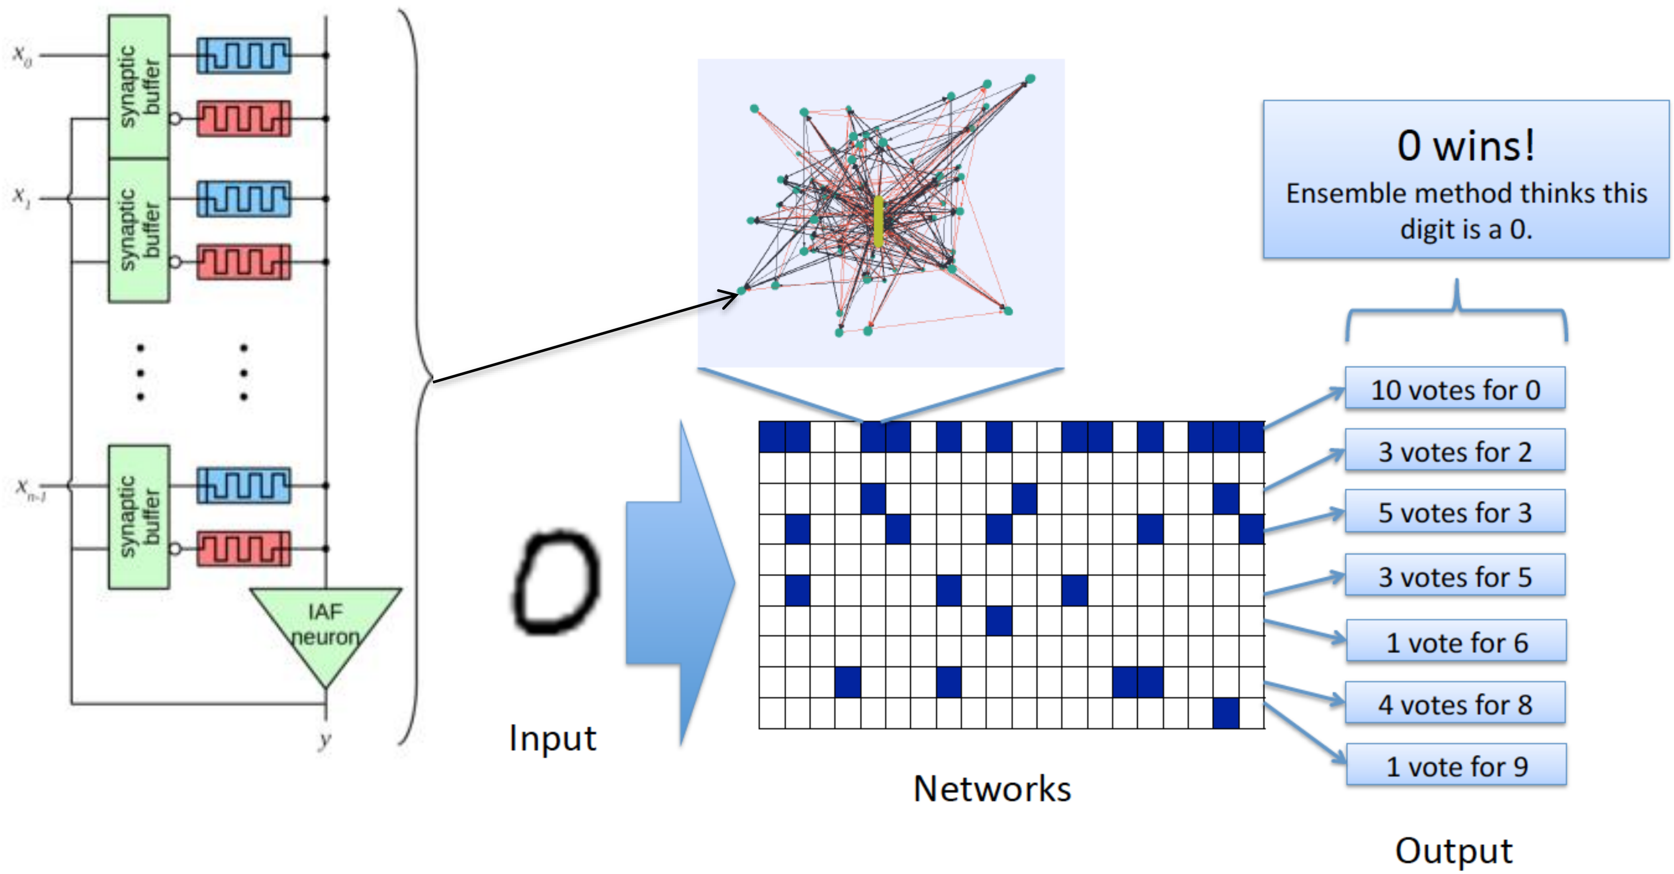

Conclusion

- Complex topologies with intra-layer connects have better classification performance than without connections
- HPC can be used to auto-tune CNN topologies
- Neuromorphic hardware has the potential to implement deep learning network in very low-power hardware
- A first step towards richer DL on novel architectures

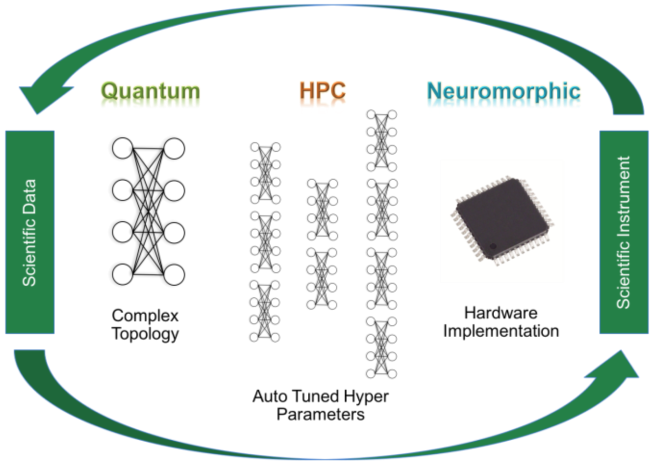

Next Step

- Evaluate the strengths and weaknesses of each approach
- Quantum: Explore more complex networks

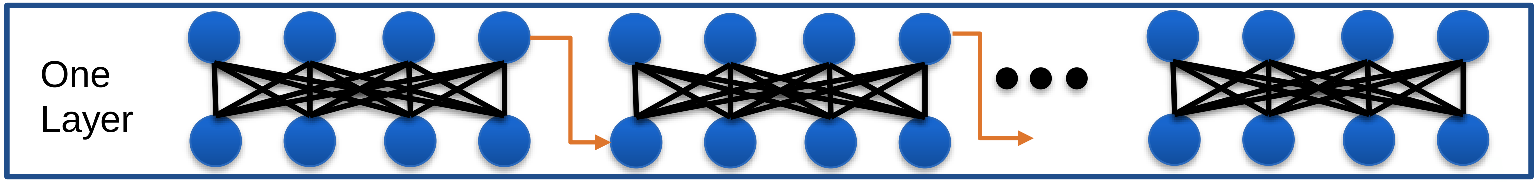

- HPC: Auto tune on Limited Boltzmann machines model
- Neuromorphic: Implement of Limited Boltzmann on FPGA version of neuromorphic hardware

OAK RIDGE National Laboratory

Team

- Deep Learning/HPC
  - Robert Patton (ORNL)
  - Steven Young (ORNL)
  - Thomas Potok (PI/ORNL)
- Quantum Computing
  - Federico Spedalieri (USC/ISI)
  - Ke-Thia Yao(USC/ISI)
  - Bob Lucas (USC/ISI)
  - Jeremy Liu (USC)
- Neuromorphic Computing
  - Garrett Rose (UT)
  - Katie Schuman (ORNL)
  - Gangotree Chakma (UT)
- Program Manager
  - Robinson Pino (DOE ASCR)

Oak Ridge National Laboratory

Questions?