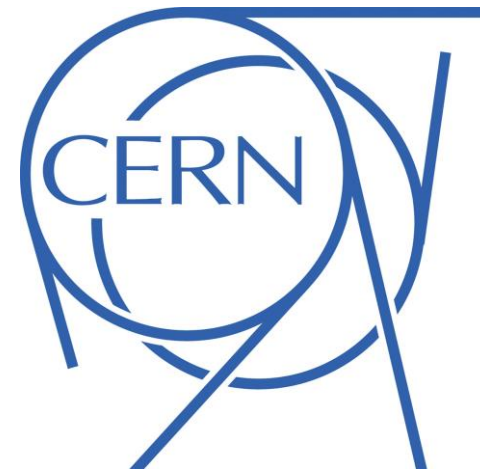


Anomaly Detection with Spiking Neural Networks on Neuromorphic Chips

- Bartłomiej Borzyszkowski (*Gdansk University of Technology, Intel Poland*)
- Eric Moreno (*Caltech*)

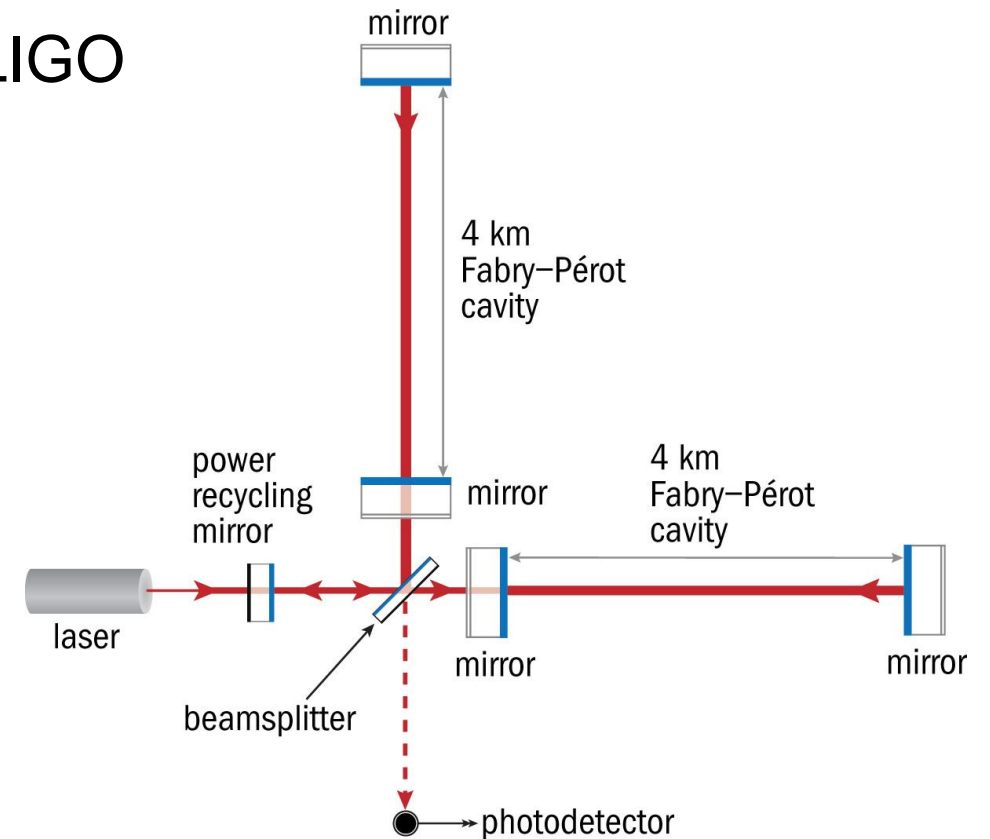
Mentors: Jean-Roch Vlimant (*Caltech*), Maurizio Pierini (*CERN*)

December 2nd, 2020; Fast Machine Learning for Science



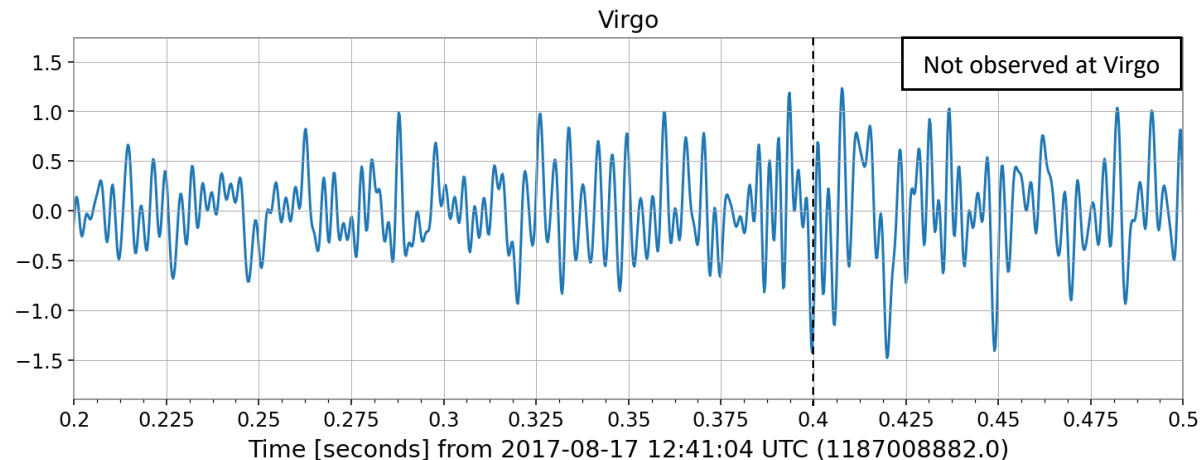
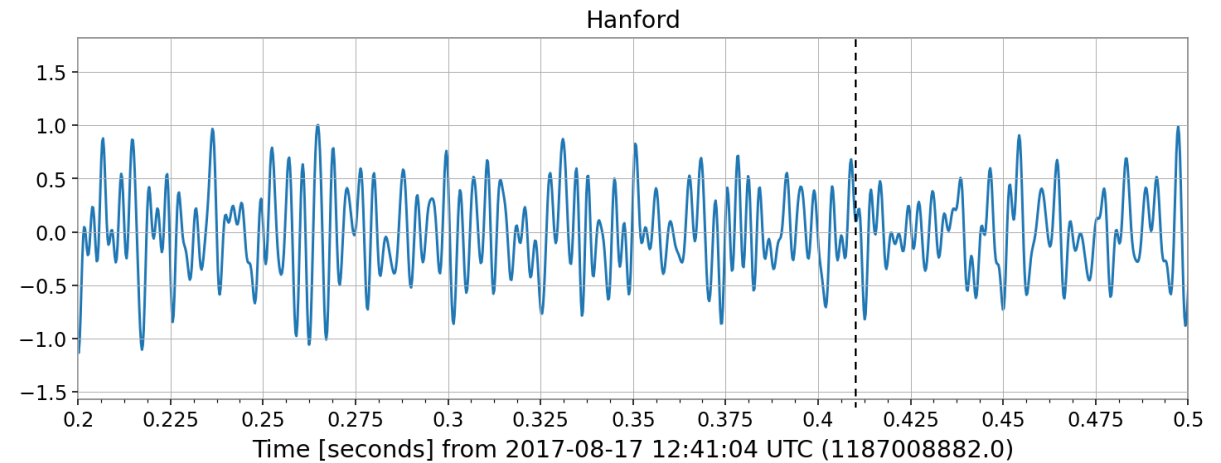
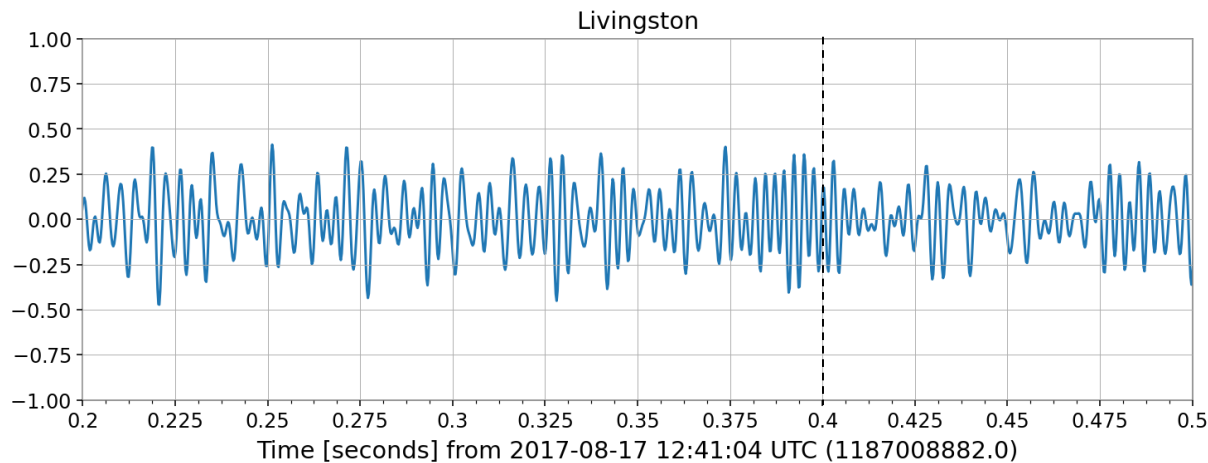
Introduction to the project

- Detection of gravitational waves (GWs) at LIGO



Produces: 1-D time-series strain

LIGO Data Generation – GW170817 example



Data is whitened and
bandpassed

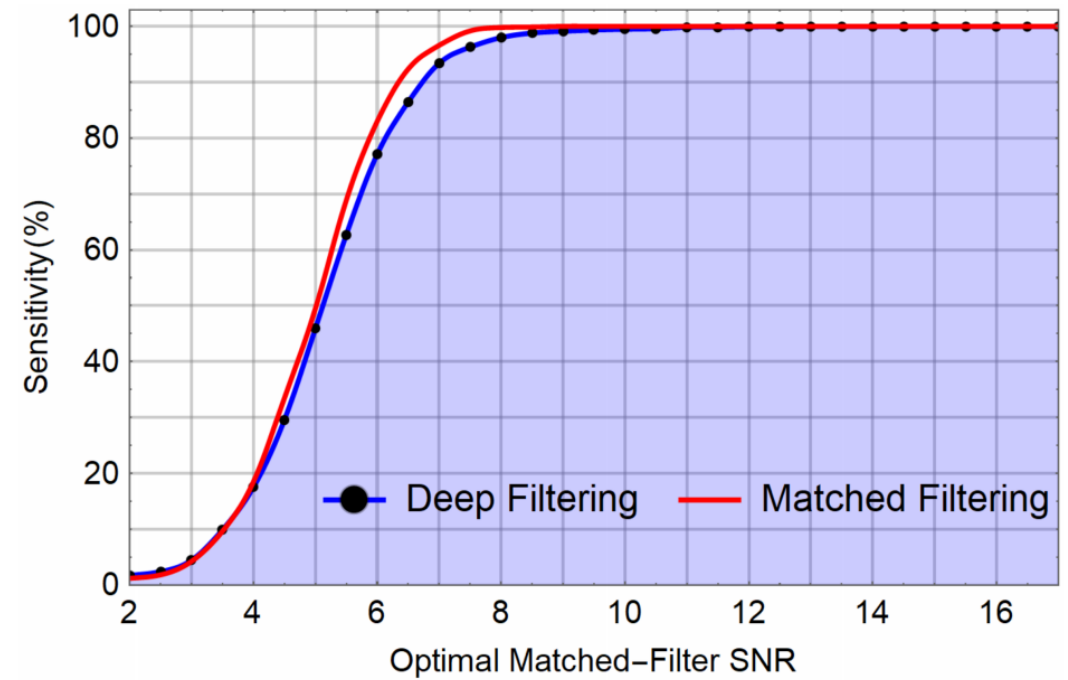
Current methods

Matched Filtering (used at LIGO)

- Compares incoming GW data to bank of simulated waveforms
- Can only identify GWs that are available in GW banks (no exotic events)

Deep Filtering (CNNs)

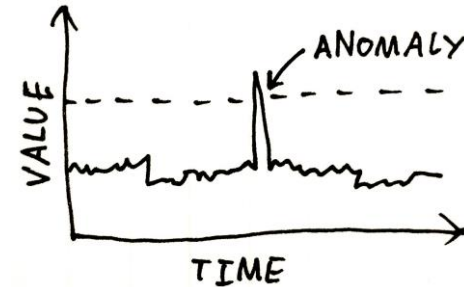
- Take time-series inputs, can determine detections and estimate parameters of events
- Still can miss events that aren't included in training set



Artificial Neural Networks

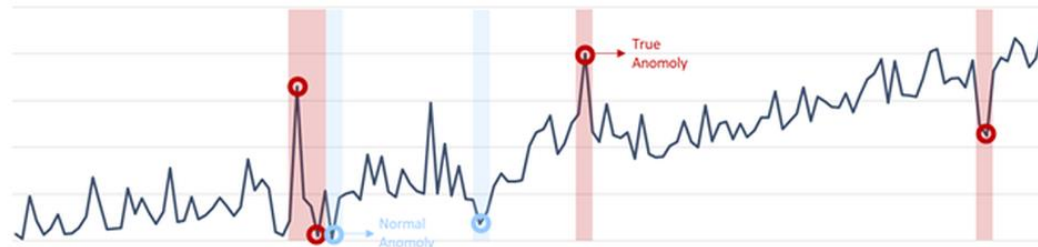
1. Unsupervised learning

- unlabeled data
- autoencoders



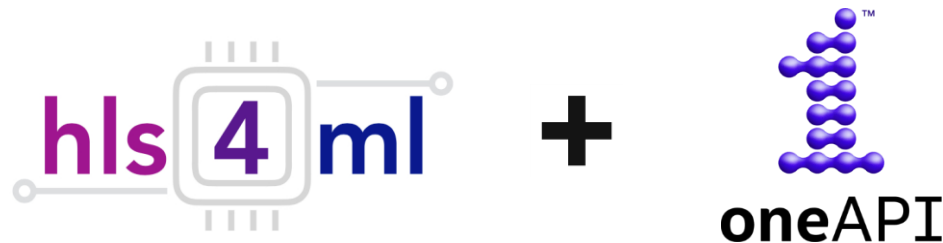
2. Supervised learning

- labels for gravitational waves
- deep neural networks

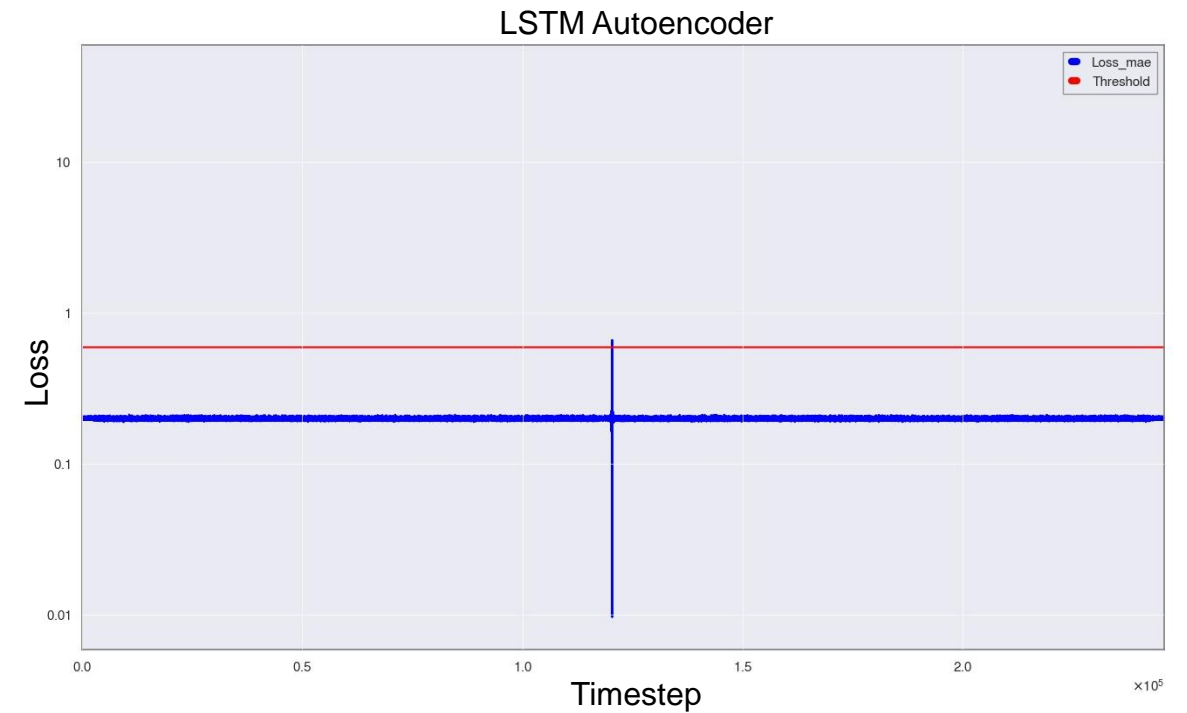
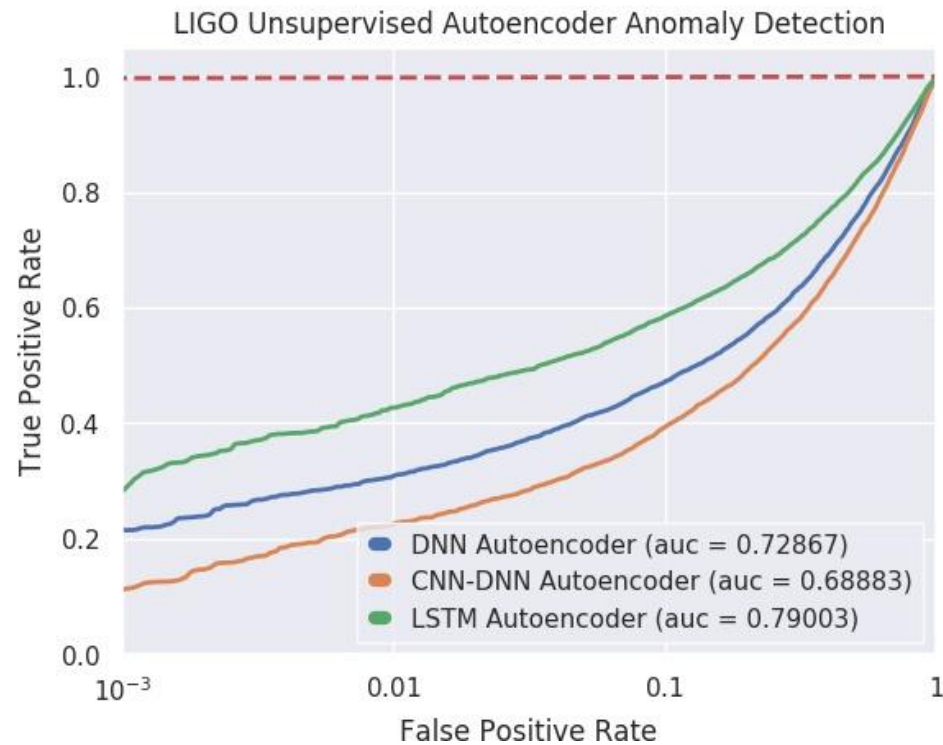


3. Fast inference

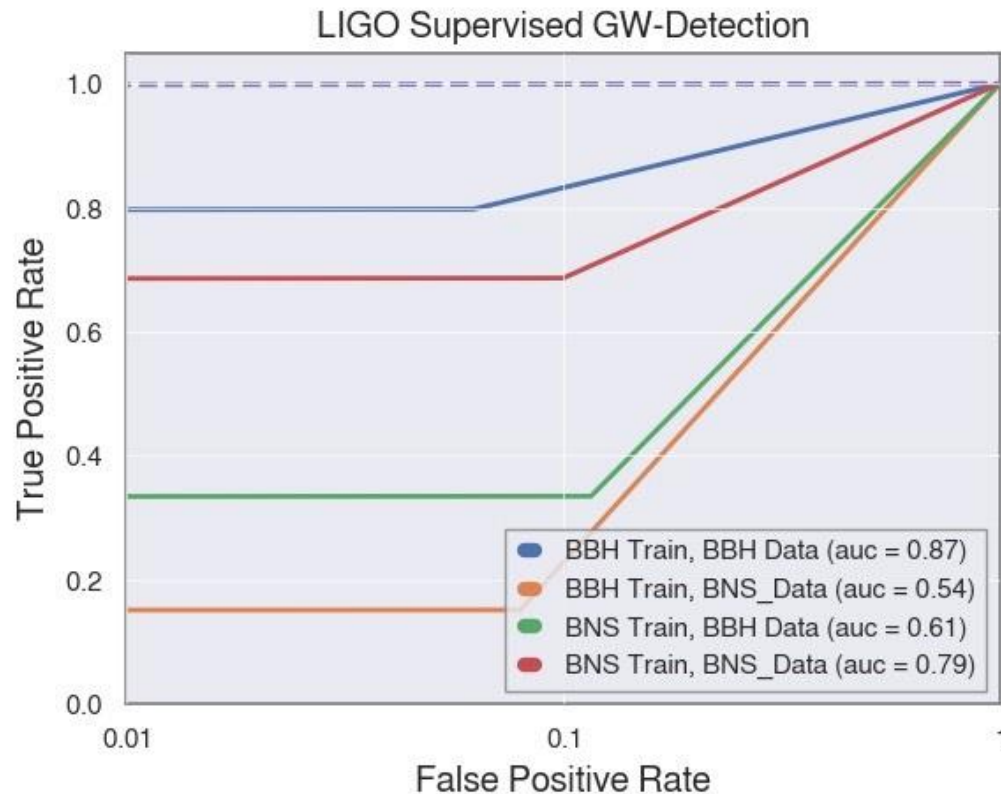
- hls4ml with oneAPI backend
- hardware acceleration



ANNs: Unsupervised learning

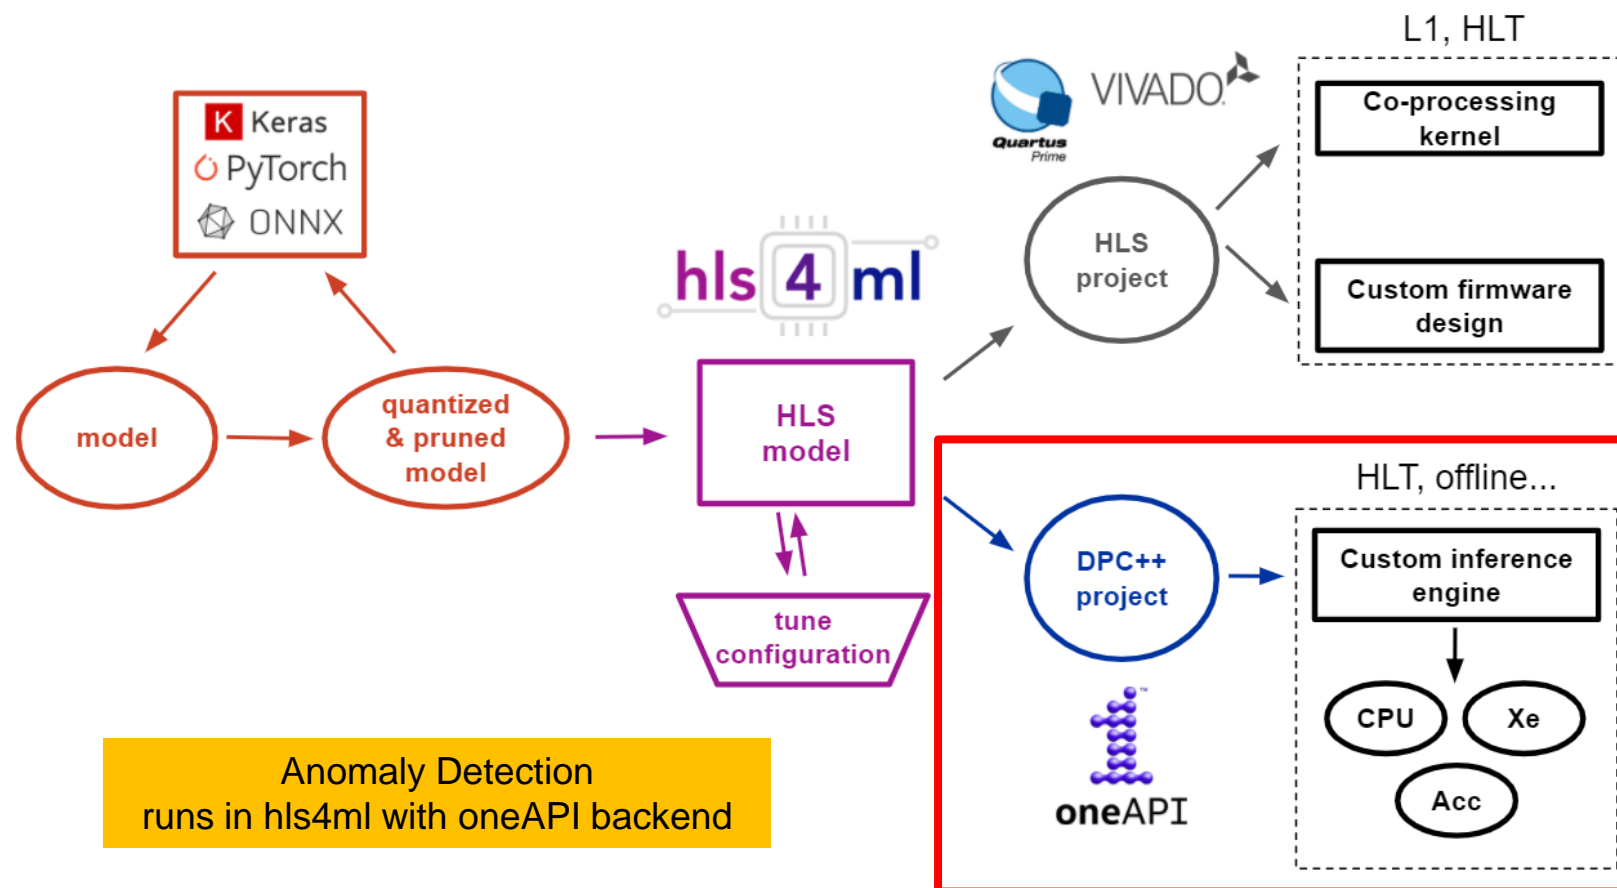


ANNs: Supervised learning



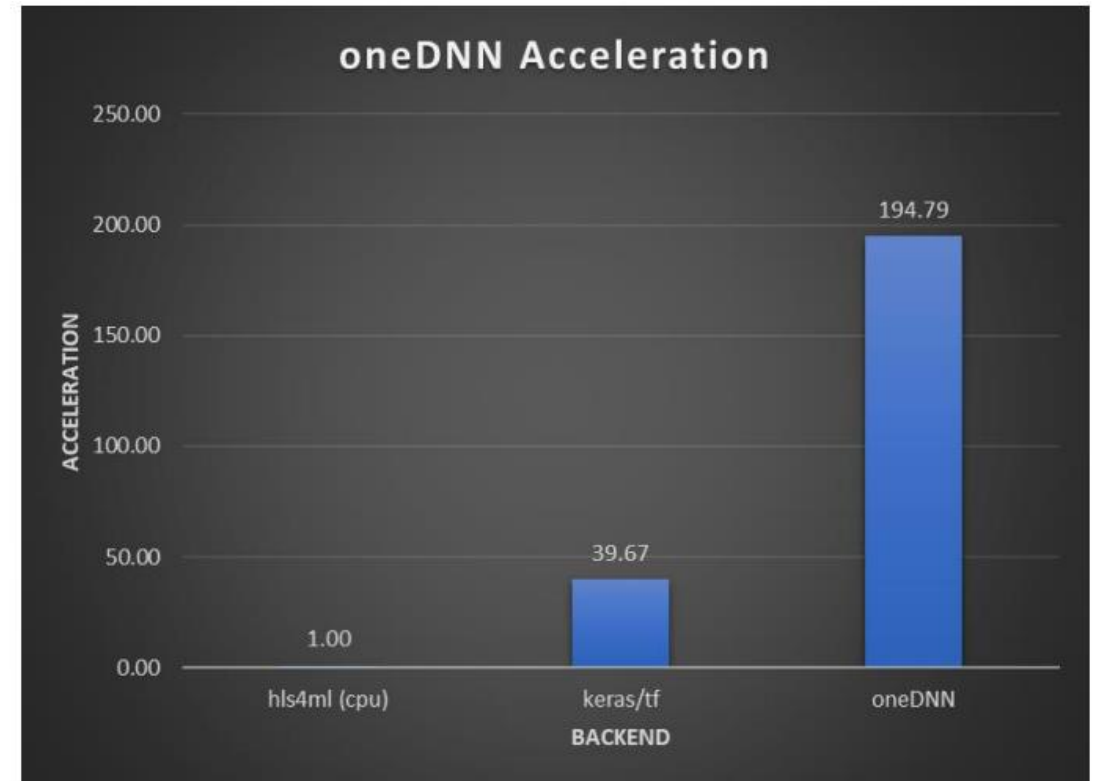
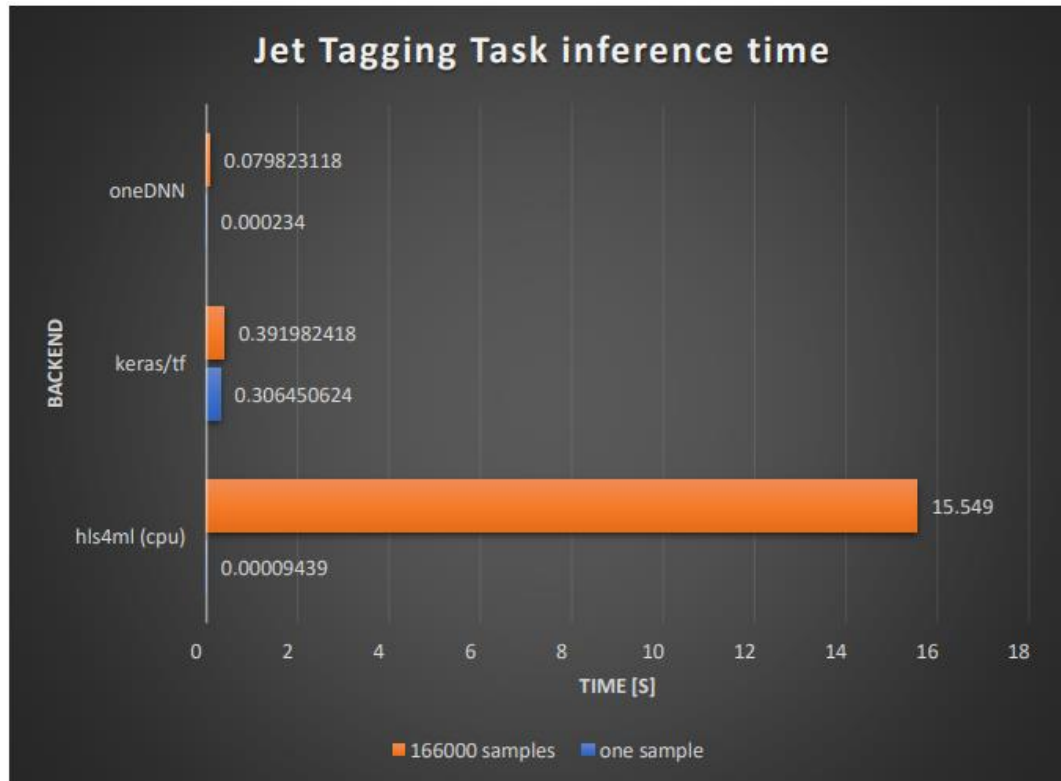
	Input	vector (size: 8192)
1	Reshape	matrix (size: 1 × 8192)
2	Convolution	matrix (size: 64 × 8177)
3	Pooling	matrix (size: 64 × 2044)
4	ReLU	matrix (size: 64 × 2044)
5	Convolution	matrix (size: 128 × 2014)
6	Pooling	matrix (size: 128 × 503)
7	ReLU	matrix (size: 128 × 503)
8	Convolution	matrix (size: 256 × 473)
9	Pooling	matrix (size: 256 × 118)
10	ReLU	matrix (size: 256 × 118)
11	Convolution	matrix (size: 512 × 56)
12	Pooling	matrix (size: 512 × 14)
13	ReLU	matrix (size: 512 × 14)
14	Flatten	vector (size: 7168)
15	Linear Layer	vector (size: 128)
16	ReLU	vector (size: 128)
17	Linear Layer	vector (size: 64)
18	ReLU	vector (size: 64)
19	Linear Layer	vector (size: 2)
	Output	vector (size: 2)

Inference engine with hls4ml and Intel® oneAPI



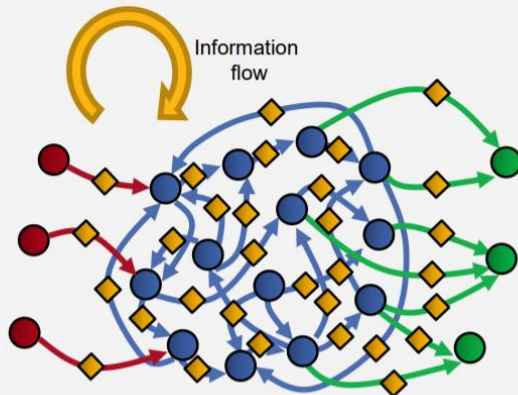
Intel® oneAPI baseline model performance

Three hidden-layer neural network on Xeon Gold 6128 3.4Ghz



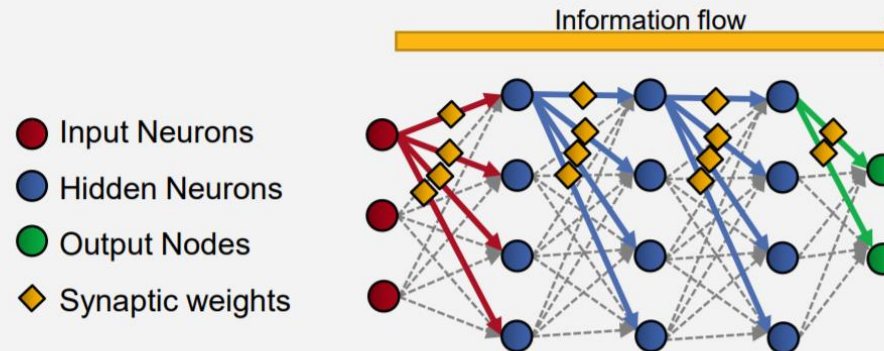
Spiking Neural Networks (SNNs)

Brain-like Neural network:



- Omni-directional signal flow
- Asynchronous pulse signals
- Information encoded in signal timing
- Inspired by biology

Deep Artificial Neural Network:



- Feed-forward sequential processing
- Use activation functions
- Information encoded in signal amplitude
- Trained with backpropagation algorithms



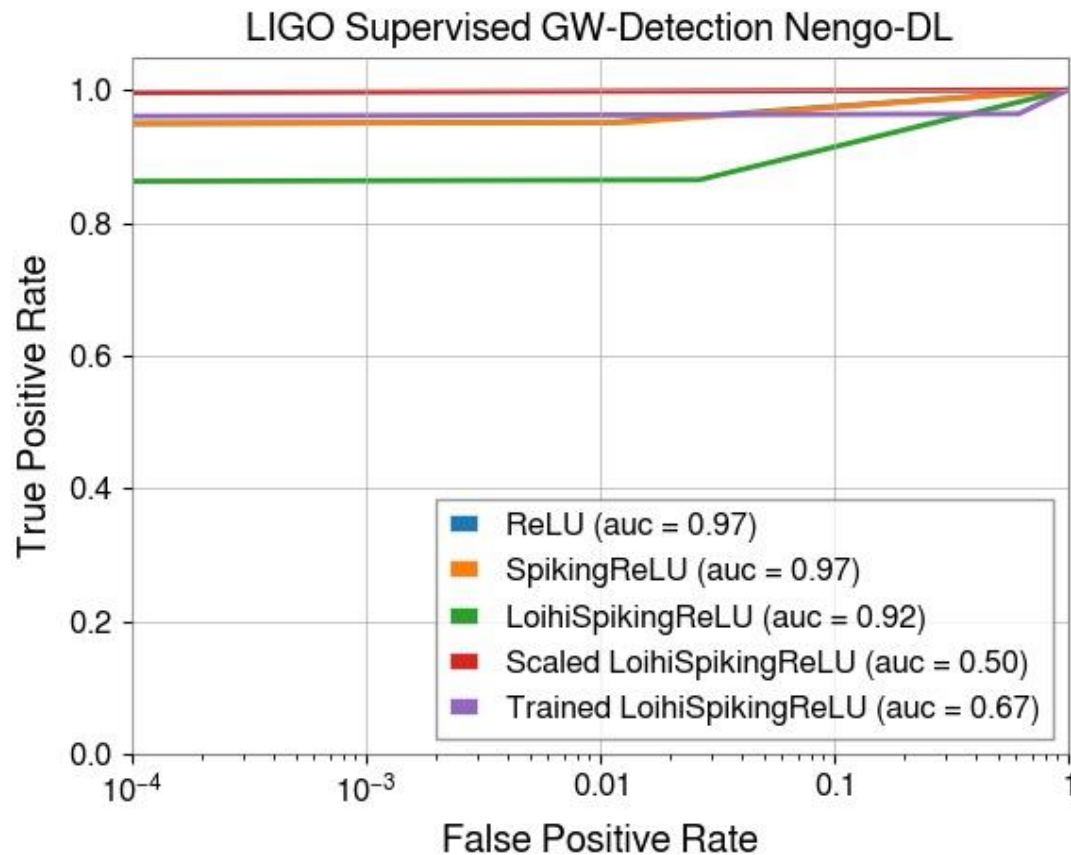
NEUROMORPHIC RESEARCH COMMUNITY

Applying inspiration from Nature for innovation
in computer architecture, algorithms, and AI

Intel Loihi – Dedicated neuromorphic chip



SNNs supervised learning - Nengo



Layer (type)	Output Shape	Param #
input (InputLayer)	[(None, 2048)]	0
reshape (Reshape)	(None, 2048, 1, 1)	0
conv2d (Conv2D)	(None, 2045, 1, 16)	64
conv2d_1 (Conv2D)	(None, 511, 1, 16)	1024
conv2d_2 (Conv2D)	(None, 127, 1, 32)	2048
conv2d_3 (Conv2D)	(None, 31, 1, 64)	8192
conv2d_4 (Conv2D)	(None, 6, 1, 128)	65536
flatten (Flatten)	(None, 768)	0
dense (Dense)	(None, 128)	98304
dense_1 (Dense)	(None, 64)	8192
output (Dense)	(None, 2)	130

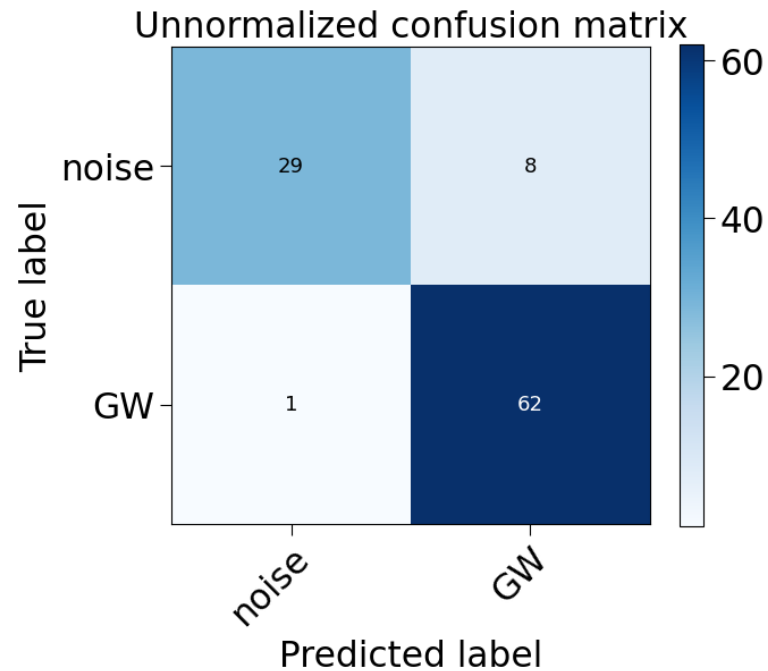
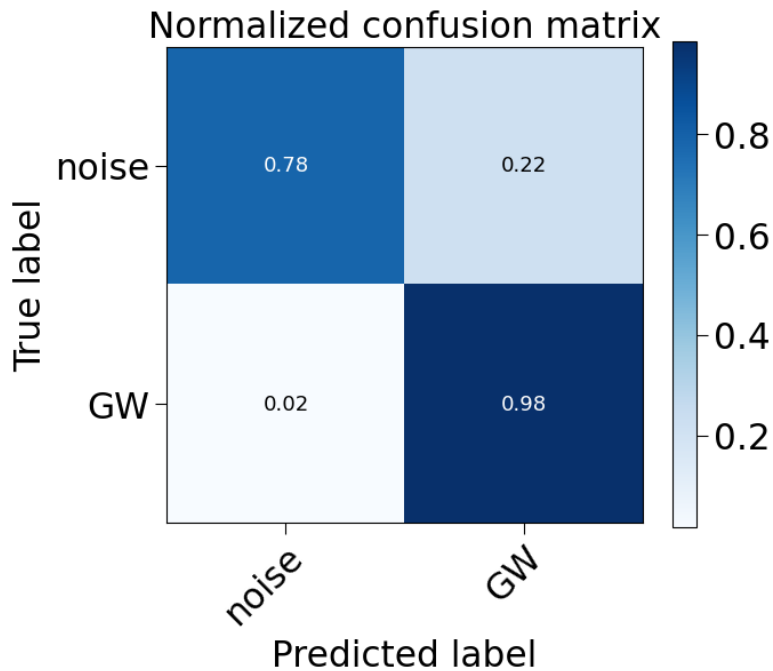
Total params: 183,490
 Trainable params: 183,490
 Non-trainable params: 0

→ Off-chip

ON-chip

Executed on Loihi chip

SNNs supervised learning - SNN-TB



Accuracy results in simulation:
ANN: 94%
SNN: 91%

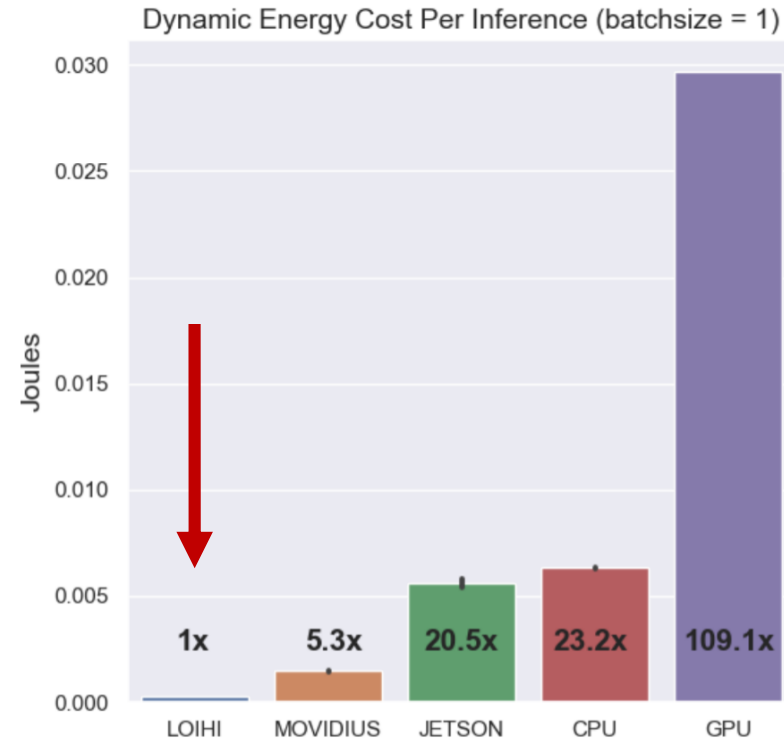
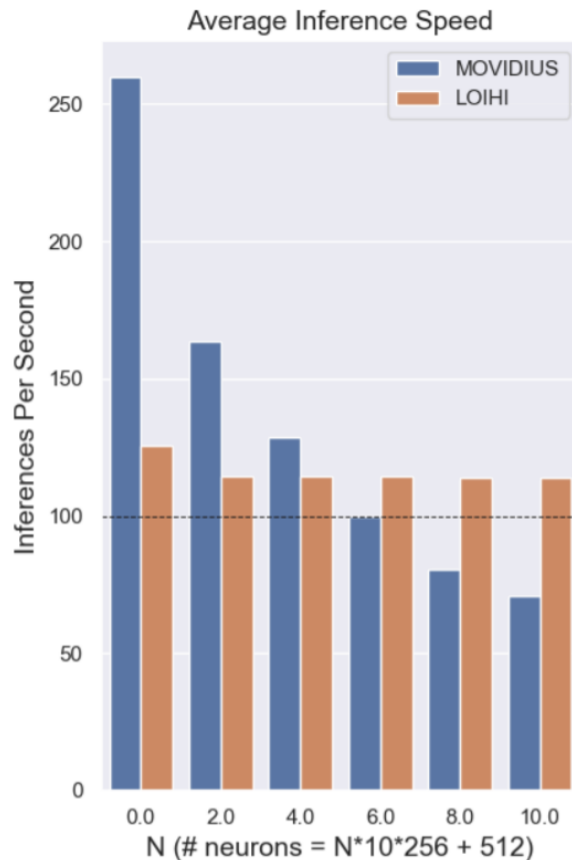
Executed on Loihi chip

Layer (type)	Output Shape	Param #
input (InputLayer)	[(10, 2025)]	0
00Reshape_2025x1x1 (Reshape)	(10, 2025, 1, 1)	0
01Conv2D_2022x1x16 (Conv2D)	(10, 2022, 1, 16)	80
02Conv2D_505x1x32 (Conv2D)	(10, 505, 1, 32)	2080
03Conv2D_126x1x64 (Conv2D)	(10, 126, 1, 64)	8256
04Conv2D_31x1x128 (Conv2D)	(10, 31, 1, 128)	32896
05Flatten_3968 (Flatten)	(10, 3968)	0
06Dense_128 (Dense)	(10, 128)	508032
07Dense_64 (Dense)	(10, 64)	8256
08Dense_2 (Dense)	(10, 2)	130
Total params: 559,730		
Trainable params: 559,730		
Non-trainable params: 0		
Evaluating parsed model on 1000 samples...		
Top-1 accuracy: 94.30%		
Top-5 accuracy: 100.00%		
Building spiking model...		
Building layer: 00Reshape_2025x1x1		
Building layer: 01Conv2D_2022x1x16		
Building layer: 02Conv2D_505x1x32		
Building layer: 03Conv2D_126x1x64		
Building layer: 04Conv2D_31x1x128		
Building layer: 05Flatten_3968		
Building layer: 06Dense_128		
Building layer: 07Dense_64		
Building layer: 08Dense_2		
Compiling spiking model...		

Selected performance benchmarks

(keyword spotting task in Nengo for DNN-SNN conversion)

Loihi provides extremely good scaling vs conventional architectures as network size grows by 50x

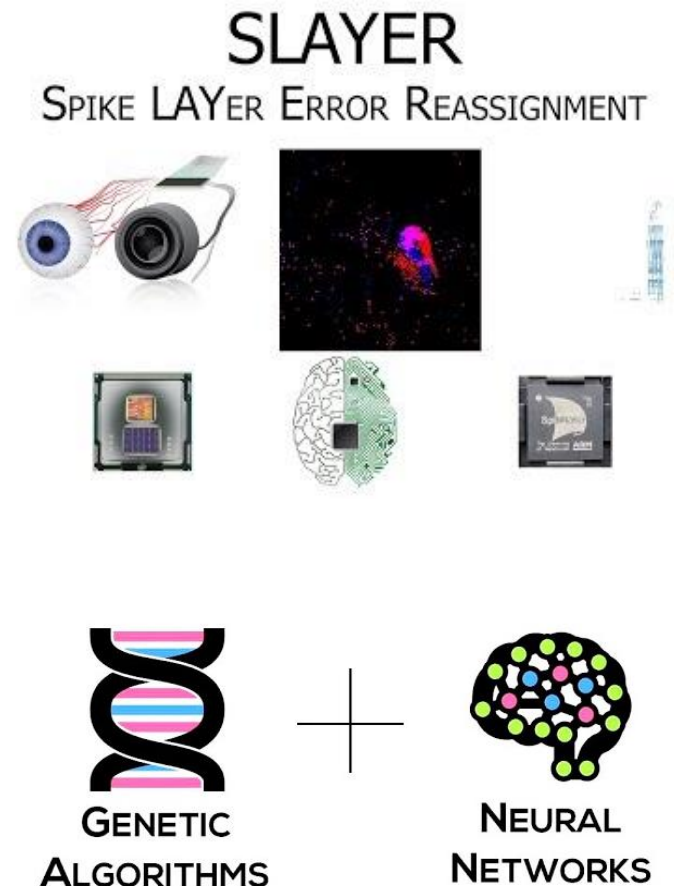


Accuracy results:
90.6% SNN
92.7% DNN

Loihi is the most energy-efficient architecture for real-time inference (batchsize=1 case)

Future opportunities for SNNs

- Unsupervised learning
(Recurrent SNNs: LSTMs, LMUs)
- Backpropagation algorithm for direct training (SLAYER)
- Evolutionary algorithms to adjust the parameters of spikes



Thank you for your attention!

Do you have any questions?

bartek_borzyszkowski@wp.pl

emoreno@caltech.edu

github.com/eric-moreno/Anomaly-Detection-Autoencoder

