fwXmachina: Anomaly detection with decision tree autoencoder on FPGA for L1 trigger

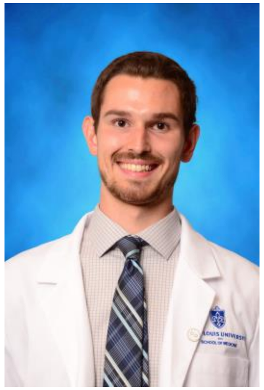

Stephen Roche*

Ben Carlson

Tae Min Hong

[2304.03836 [hep-ex]](2304.03836 [hep-ex])

Outline

- Autoencoder training
- Results
- Comparison to neural network

Fast Machine Learning for Science

September 25, 2023

https://indico.cern.ch/e/1283970/contributions/5554363/

Autoencoders for anomaly detection

Stephen Roche

One anomaly detection method widely used is autoencoders

- Initially devised for data compression
- Use some method (often neural networks) to compress data into a latent space
- A second neural network can be used to take latent space → reconstructed object
- If input looks like training data (background), good reconstruction. If not (anomaly), bad reconstruction

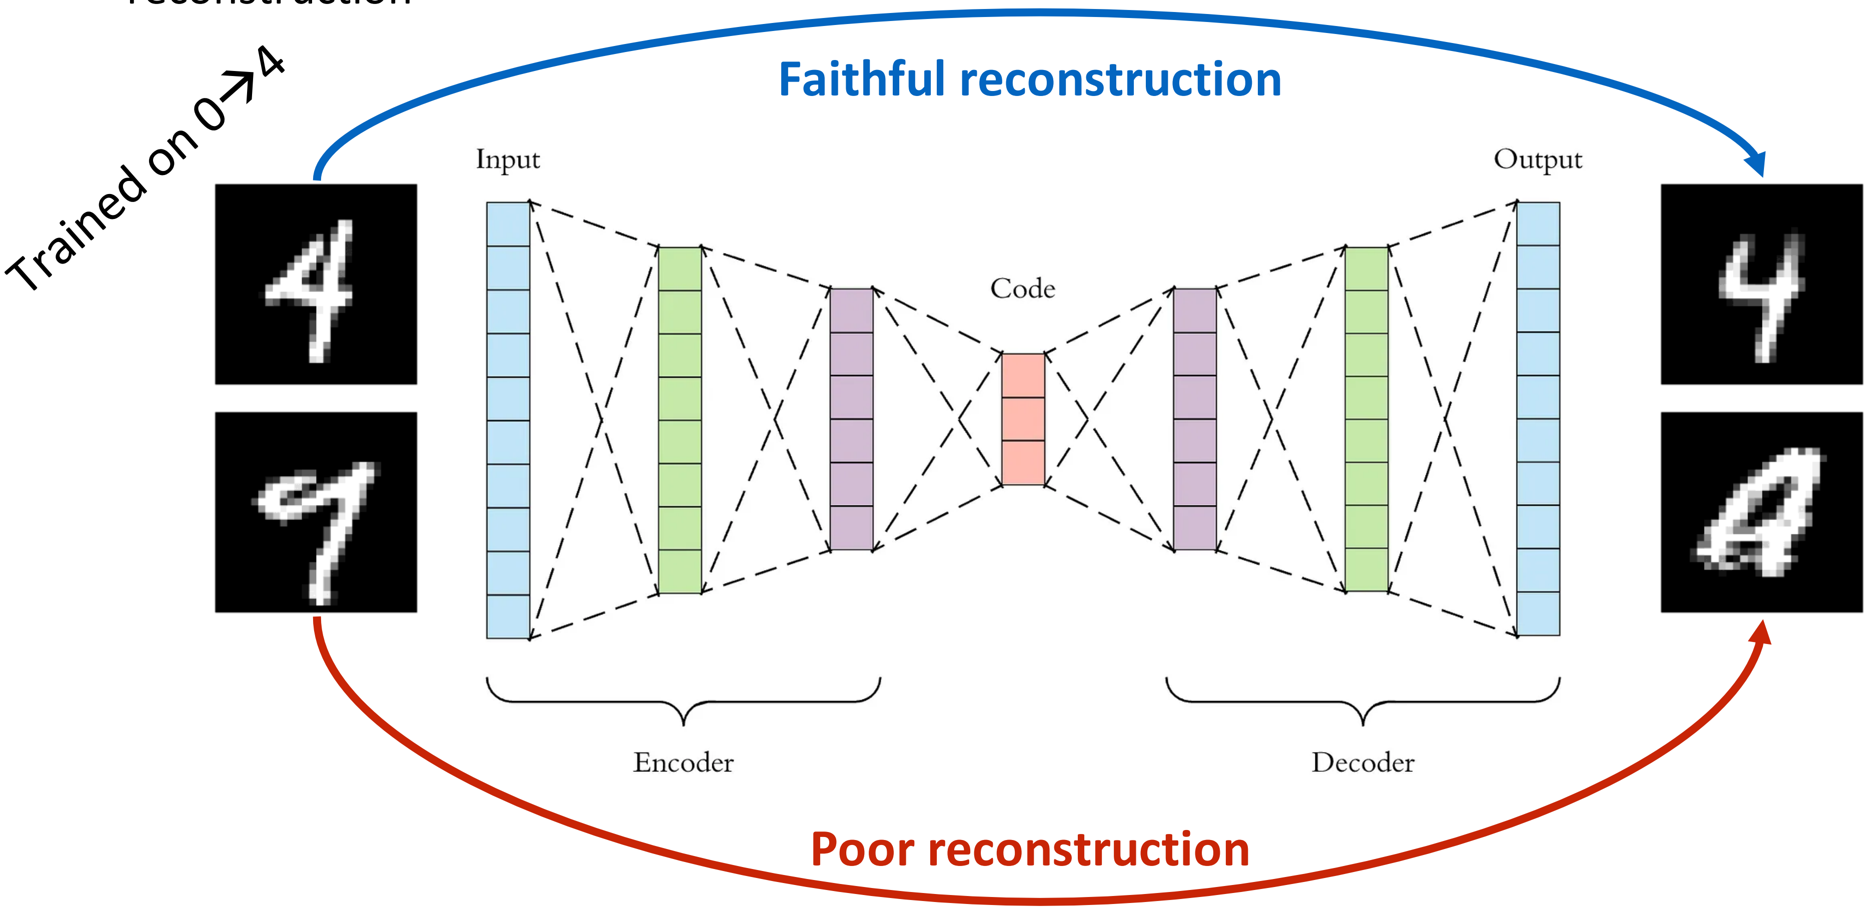

Figure: https://towardsdatascience.com/applied-deep-learning-part-3-autoencoders-1c083af4d798

Decision tree autoencoders

Stephen Roche

We developed an algorithm that trains **decision trees** (rather than NNs) as autoencoder

- Can be evaluated on FPGA with fwXmachina framework (see slides from earlier overview talk by Tae Min Hong https://indico.cern.ch/event/1283970/contributions/5554356/)
- Based on density of background points in parameter space
- The MNIST digits on the last slides were evaluated using decision tree AE
- See paper for training algorithm details

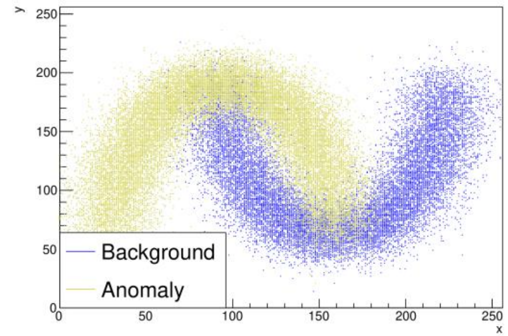

Consider 2D toy dataset

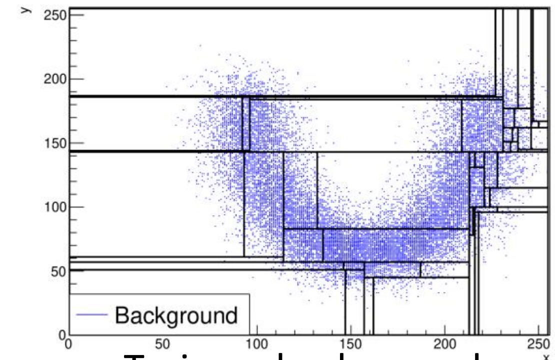

Train on background

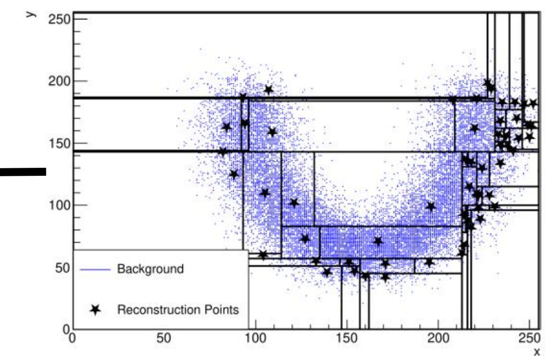

Find reconstruction points

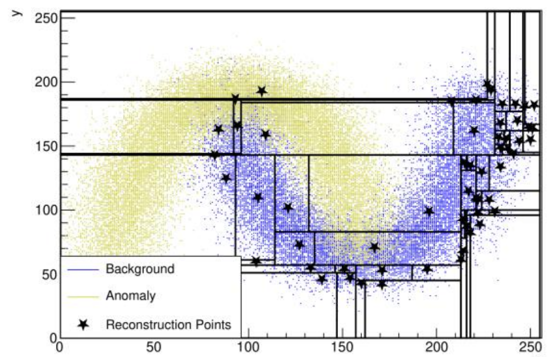

Examine anomaly scores

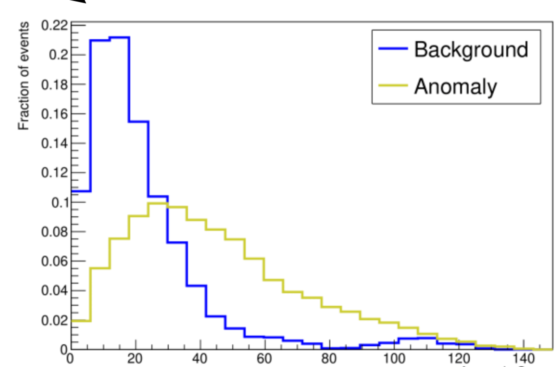

Evaluate on background & anomaly

Results: Exotic Higgs decay

Tested on dataset: $e^+e^-\mu^+\mu^-$ background vs two BSM $H \rightarrow aa \rightarrow e^+e^-\mu^+\mu^-$ signals with different Higgs and pseudoscalar (a) masses

- Only included events that wouldn't pass single lepton trigger
- Trained on background

Parameter	Value
Variables	3 (m_{ee}, $m_{\mu\mu}$, $m_{ee\mu\mu}$)
Configuration	40 trees, depth 5
Clock speed	320 MHz
Latency	8 ticks (25 ns)
Interval	1 tick (3.125 ns)
FF	0.4%
LUT	2.6%
DSP	0.04%
BRAM	0

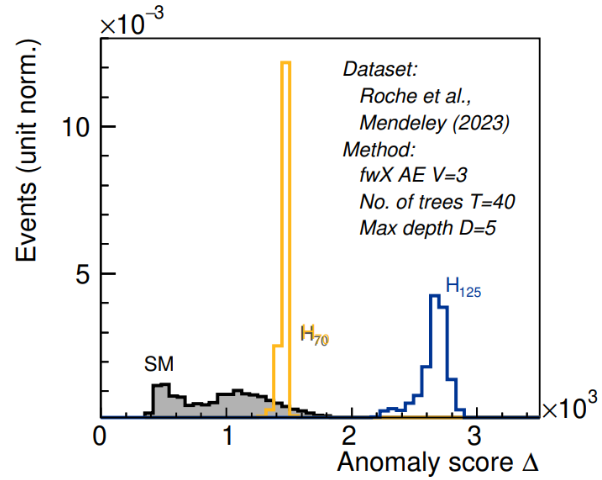

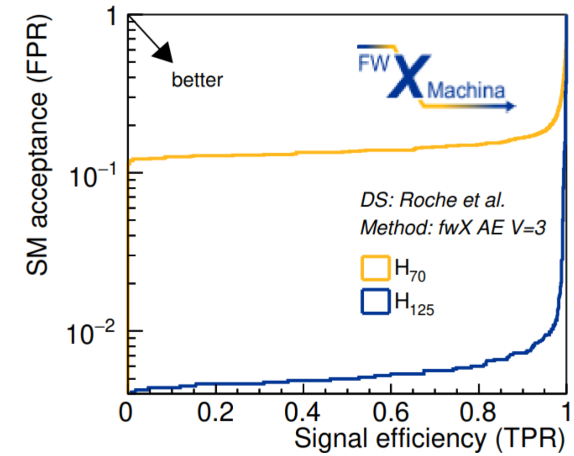

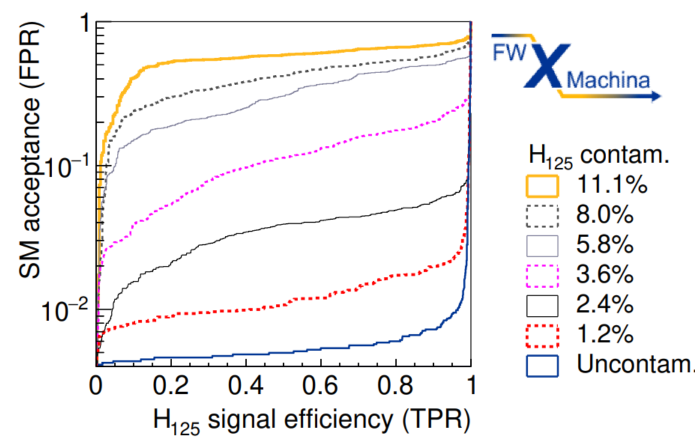

Can train with some signal contaminating the training set without significant decrease in performance

- Possibility to **train on data** rather than simulated samples

Comparison to NN autoencoder

Stephen Roche

Compare our tool to public results from hls4ml: [2108.03986]

- Both perform very well on problem using 56 variables, 4-vectors of physics objects for several signals vs SM cocktail background
- fwX has lower latency, hls4ml has less LUT usage on this problem

Signal	Area under ROC curve	
Dataset: [2107.02157]	fwX	hls4ml
LQ → bτ	**0.93**	0.92
A → 4l	0.93	**0.94**
h→ ττ	**0.85**	0.81
h^+ → τν	0.94	0.94

Parameter	fwX	hls4ml
Variables	56	56
Configuration	30 trees, depth 4	DNN VAE PTQ
Bit precision	8	8
Clock speed	200 MHz	200 MHz
Latency	30 ns	80 ns
Interval	5 ns	5 ns
FF	0.6%	0.5%
LUT	9.2%	3.0%
DSP	0.8%	1%
BRAM	0%	0.3%

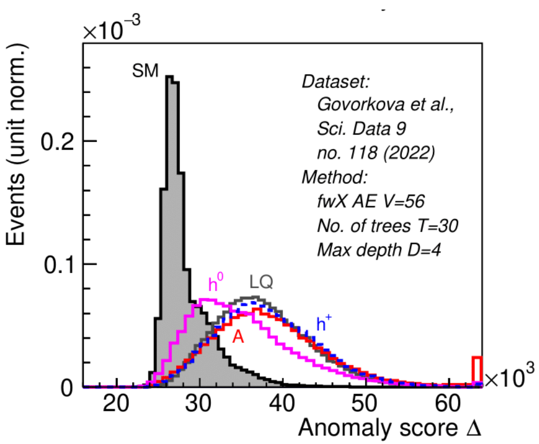

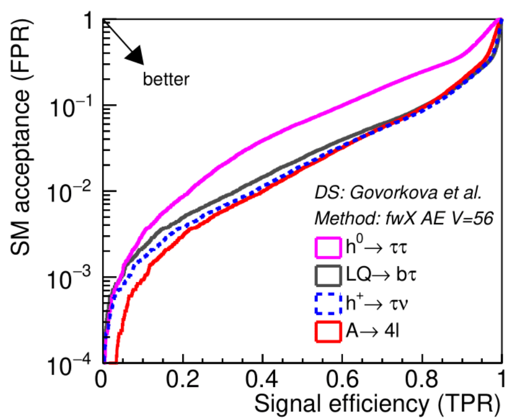

Discussion & conclusions

We have developed a novel algorithm for training decision trees as autoencoders for anomaly detection

- Allows for interpretable anomaly detection
- Can be implemented on FPGA for ultra-low latency evaluation with fwX platform
- Different tools are available for FPGA-based anomaly detection, each with strengths and weaknesses

Anomaly detection at L1

- Signal-agnostic anomaly detection can enable the L1 trigger to save BSM events that would otherwise be discarded
- Important to ensure we are not discarding new physics!

Questions?

stephen.roche@health.slu.edu

Anomaly detection in HEP

- Anomaly detection (AD) is a topic in HEP of much current interest
- Lots of recent papers on methods (HEP ML Living Review); ATLAS analysis recently performed using anomaly detection [2307.01612]
- Can't analyze events you aren't saving! We want to apply **AD methods at L1 trigger** to ensure we're not discarding new physics

fwXmachina

- fwX framework evaluates BDTs on FPGA
- Classification [2104.03408], regression [2207.05602], now autoencoder [2304.03836] (this talk)
- See slides from earlier overview talk by Tae Min Hong https://indico.cern.ch/event/1283970/contributions/5554356/

Backup: LHC anomaly detection dataset

Stephen Roche

We test our method on the hls4ml dataset [2107.02157]

- Background: cocktail of SM processes including W → ν l, Z → l l, multijet, and ttbar
- Signal: 4 different BSM decays
- Variables are p_T, η, ϕ of the 4 leading muons, 4 leading electrons, 10 leading jets, and MET
- Only events with at least one lepton > 23 GeV are included