PURM
Penn Undergraduate Research Mentoring Program

A Machine Learning-Based Approach to Real-Time Anomaly Detection in pp Collisions at the LHC

Anirudh Bharadwaj[1], Dylan Rankin[2]
University of Pennsylvania (SEAS '26)[1], University of Pennsylvania Department of Physics & Astronomy[2]

Background

- Large Hadron Collider accelerates and collides beams of protons together
  - Result: hundreds of millions of proton-proton collisions
- Vast majority of collisions produce jets of low-energy hadrons
  - In some theories of new physics, vanishingly small amount of collisions do produce anomalous phenomena.
- Detecting signal events in real time necessitates model that is:
  - Capable of filtering out the vast majority of data while ensuring atypical information is kept
  - Simple enough to meet hardware requirements

Methodology

- Model efficacy was determined by testing on 4 datasets, each comprising one of the following signal events:
  1. Zvvhbb events: 2 bottom quarks, 2 neutrinos
  2. Ztt events: 2 tau's
  3. Ttbar events: One top quark, one bottom quark
  4. Vbfhhbbbb events: 4 bottom quakers
- 4 different models were tested: a supervised model, an autoencoder, and two variational autoencoders.

Supervised Model

- Provided already labeled data (signal events labeled 1, background events labeled 0)
- Trained on a mix of background (non-anomalous) events and single signal event.
- Loss Function: binary cross-entropy

Autoencoder Model

- Unsupervised neural networks that reconstruct the input data
- Custom loss function: mean squared error (MSE) b/w input and output if output > input, MSE * 0.1 otherwise

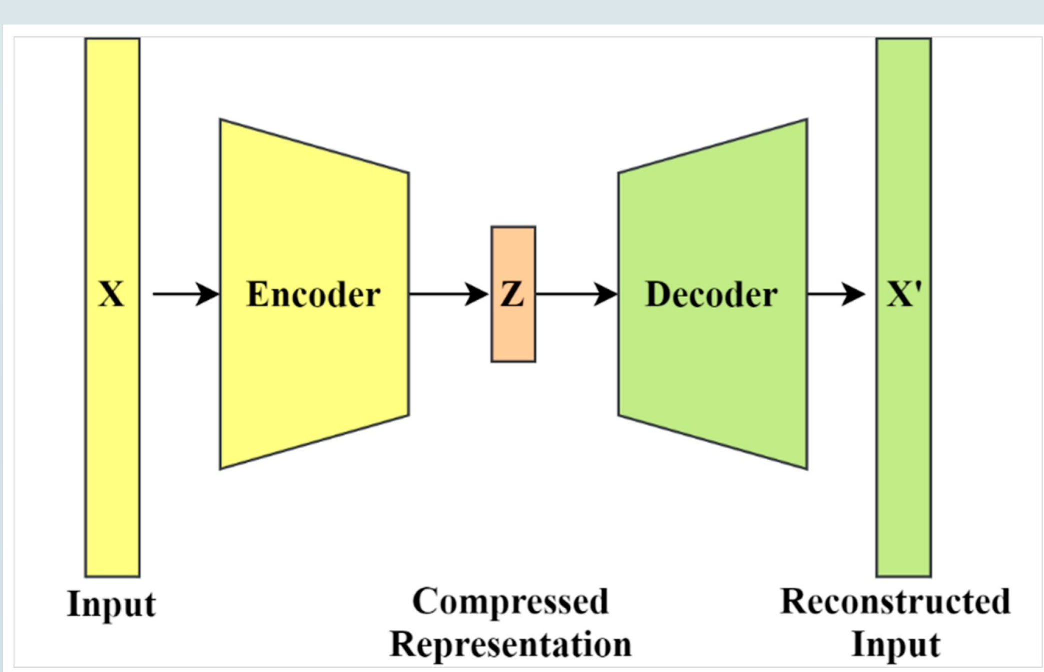

Variational Autoencoder Models

- Similar in structure to Autoencoders
- Latent space is normalized distribution → points randomly sampled
- Loss function consists of a reconstruction loss and a latent space loss
  - Two possible metrics for reconstruction loss: mean squared error (MSE) and earth mover's distance (EMD)
  - One VAE constructed using each
  - Separate model trained and utilized to estimate EMD

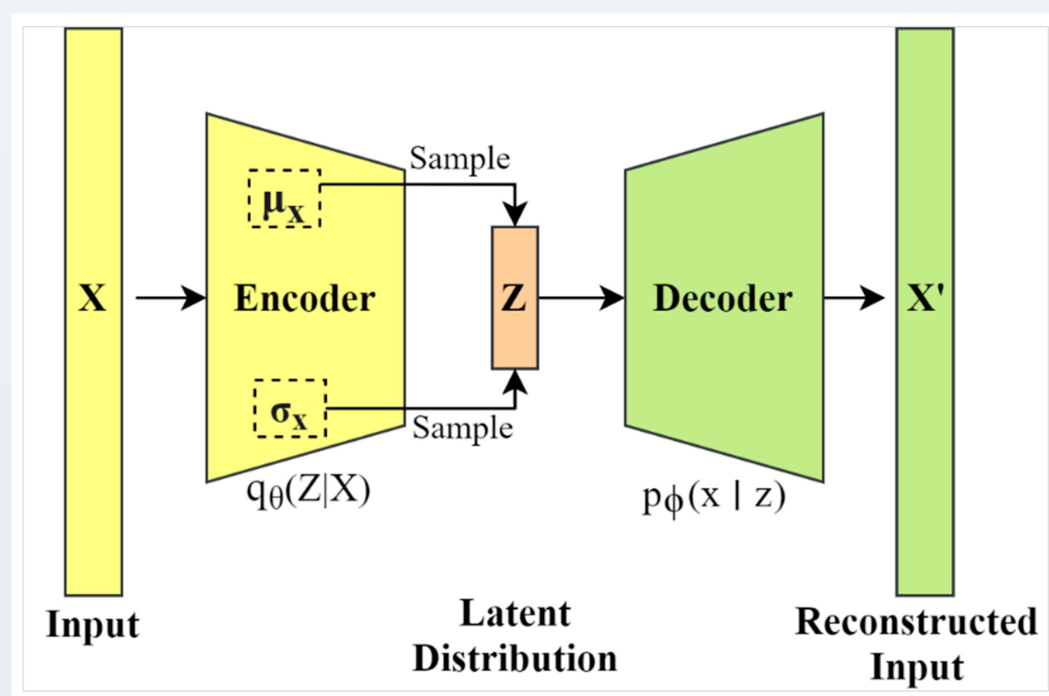

Results

- Relative utility of each model → ROC Curves & Efficiency Plots

ROC Curves

- Plots false positive rate against the true positive rate at multiple thresholds
- Determine model that would maximize true positive rate at a false positive rate of 10^-4

Efficiency Plots

- Plots ratio of events labeled anomalous by model to total events, binned by MET (missing energy in each event)
- Key:
  1. Supervised Model
  2. Autoencoder Model
  3. Variational Autoencoder (MSE)
  4. Variational Autoencoder (EMD)

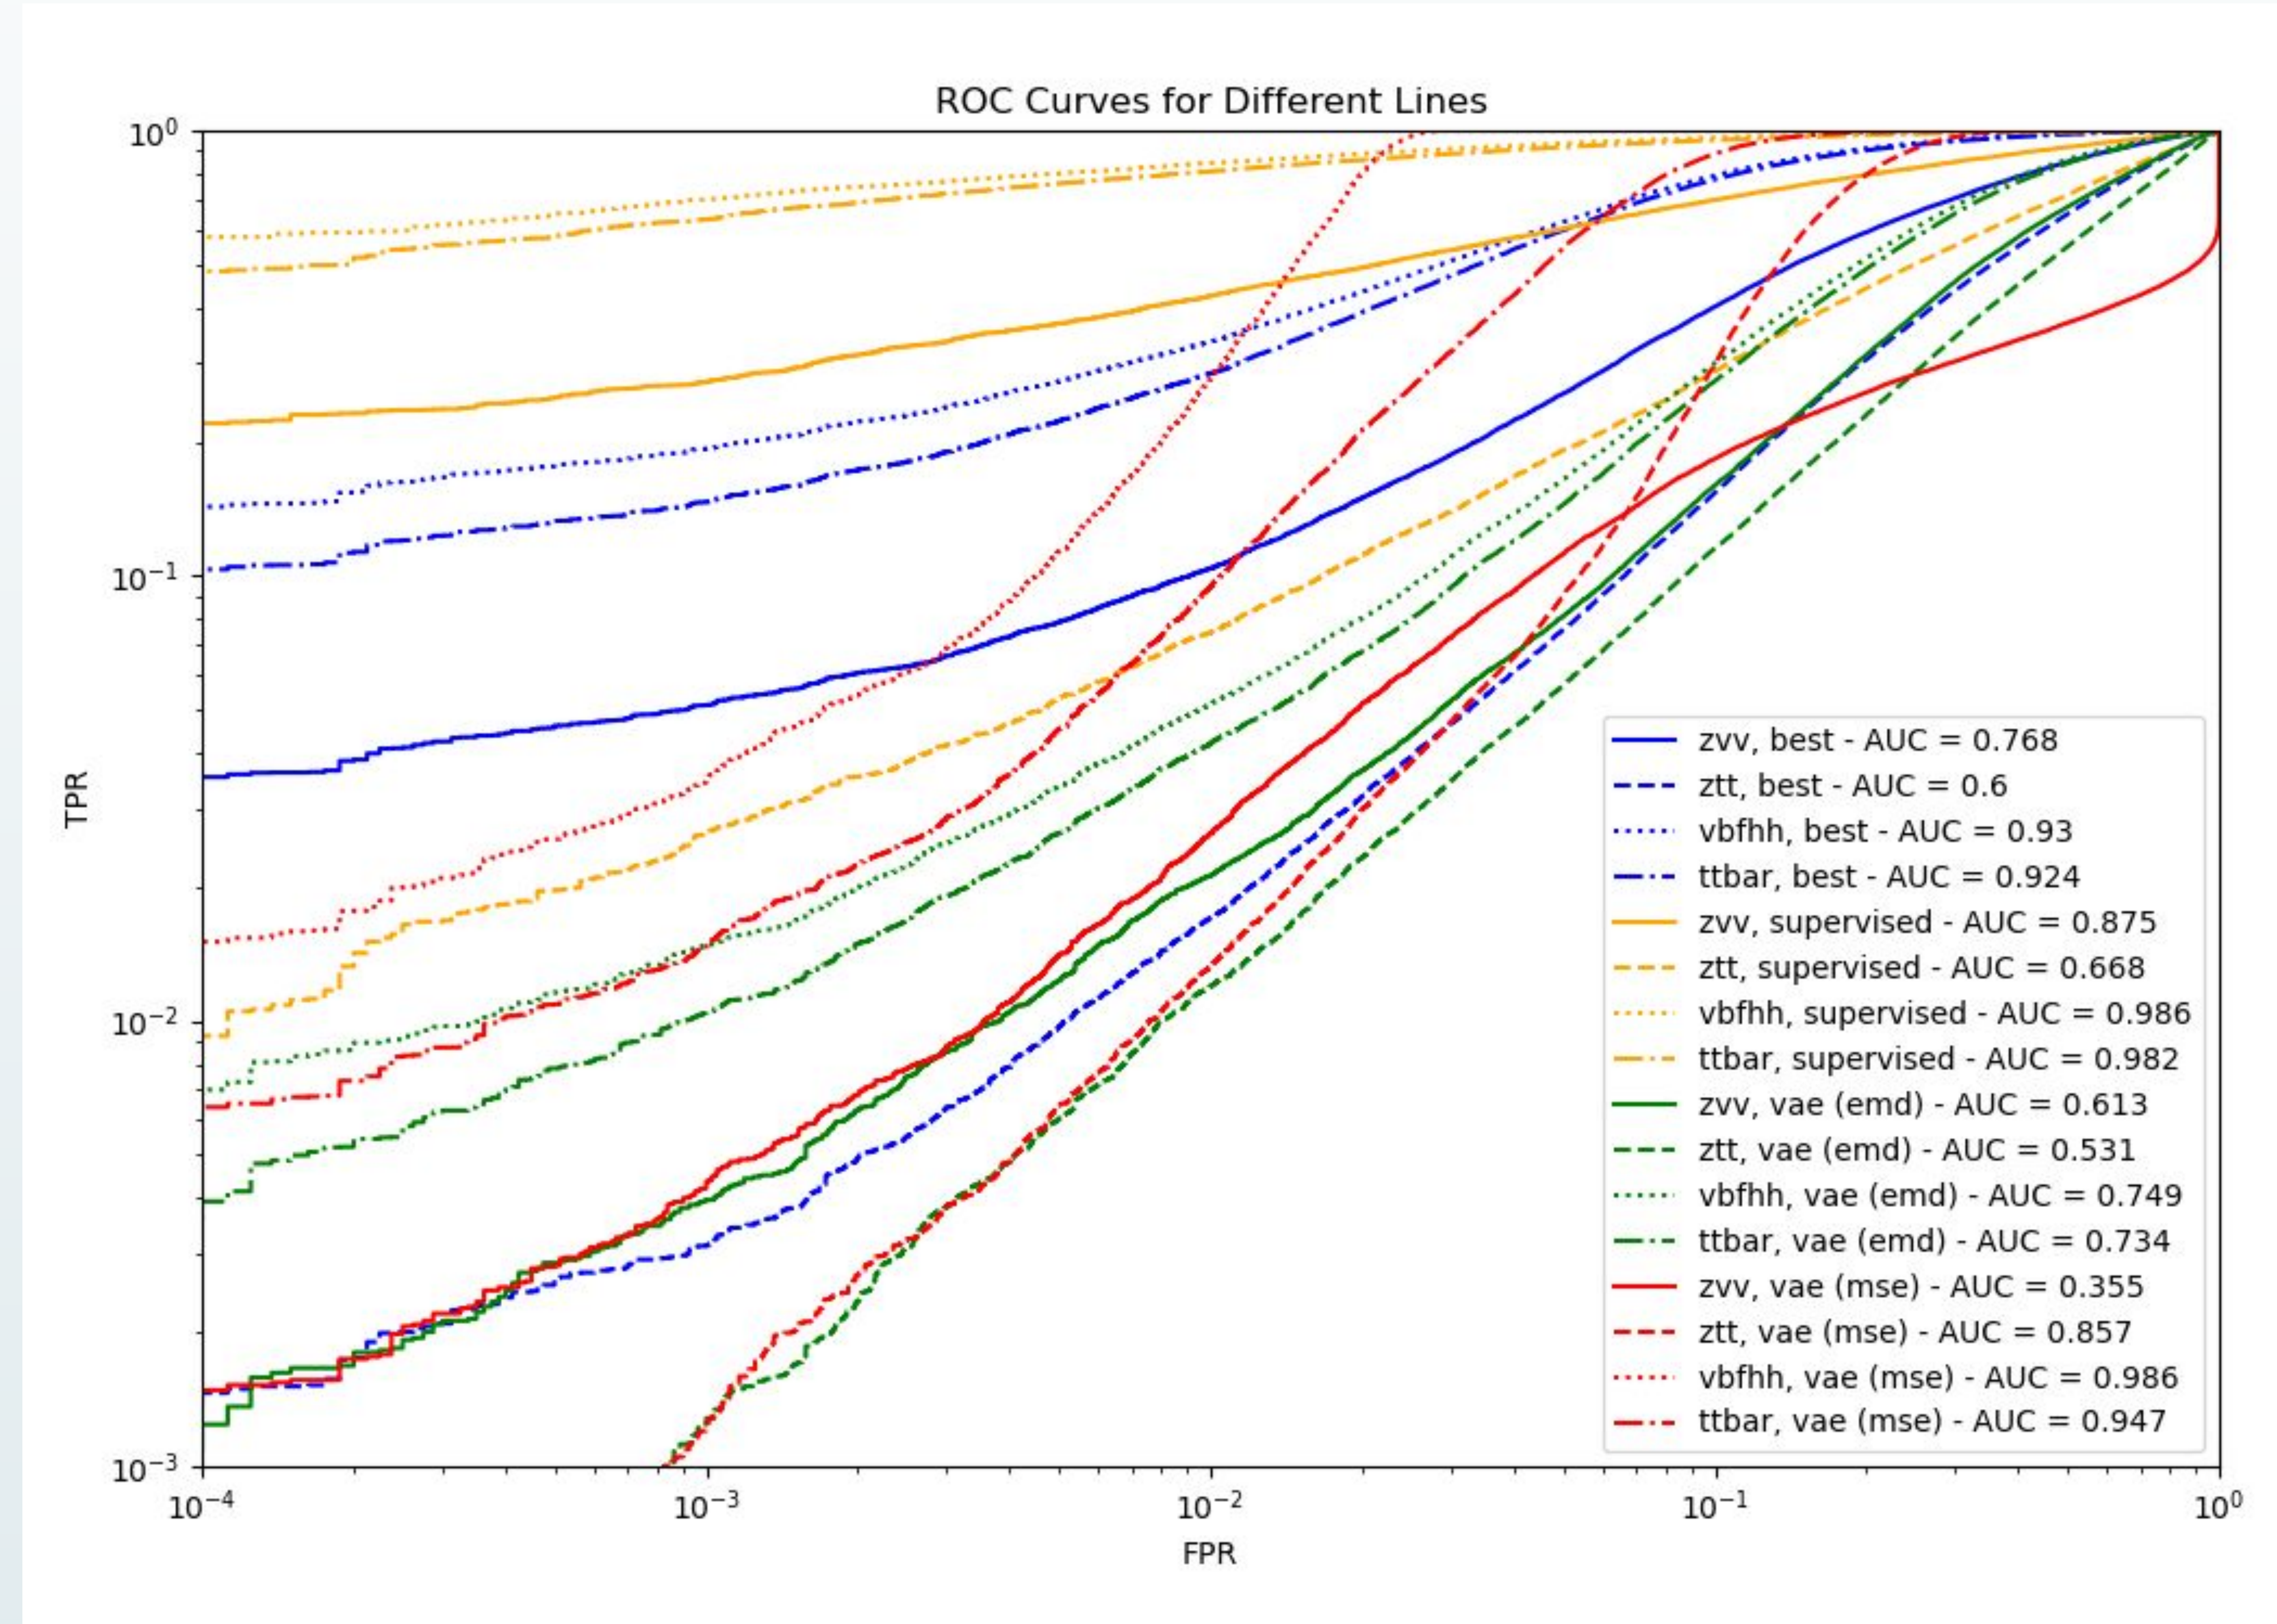

Efficiency Plots (cont.)

1.
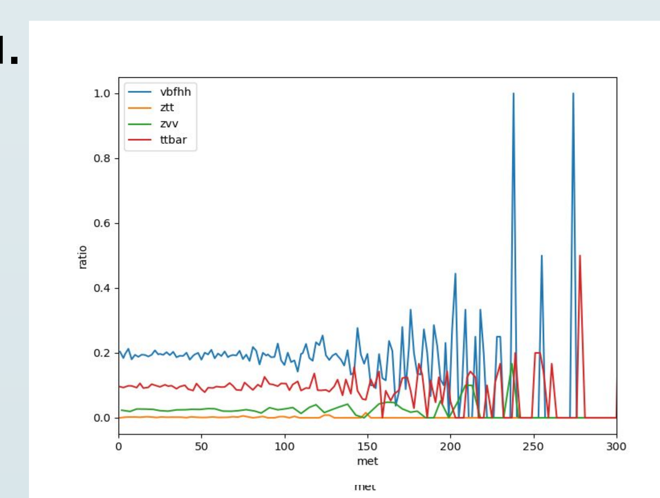

2.
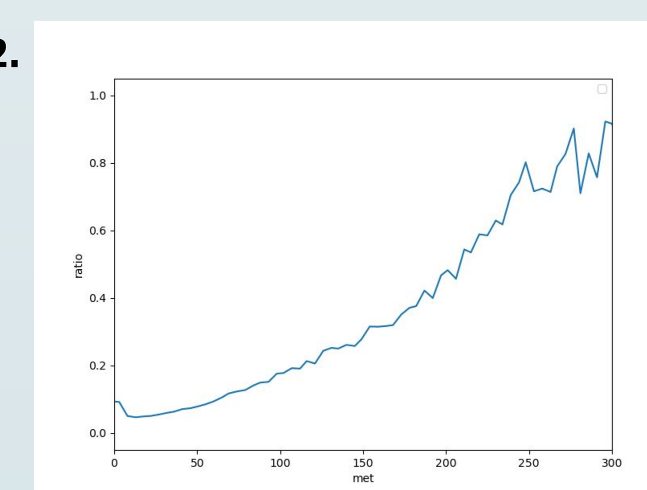

3.
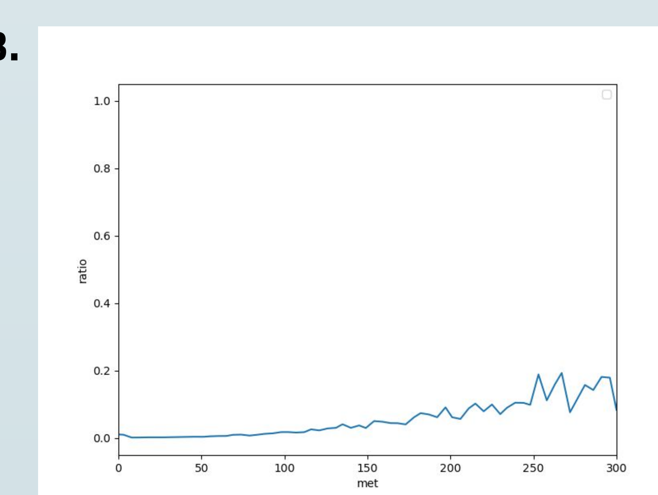

4.
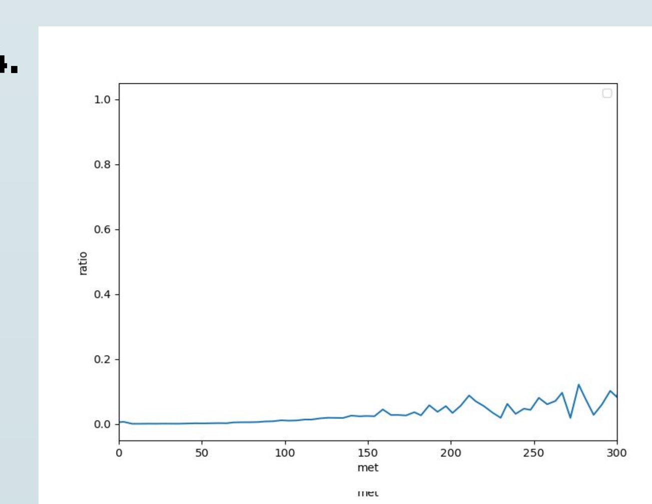

Conclusions

- Supervised model provides best tradeoff b/w TPR & efficiency
- Autoencoder appears to provide better tradeoff than either VAE
  - Worse efficiency than either VAE made up by large disparity in TPR

Future Directions

- Consider models based on additional autoencoder architectures
- Look at efficiency of each model w/ respect to alternate metrics (ex. total energy in each event)