Anomaly Detection in the LZ Dark Matter Experiment using Autoencoders

Tyler Anderson

Stanford University, SLAC National Accelerator Laboratory

Motivation

Characterizing backgrounds is an extremely important task for any dark matter search, and almost every new detector encounters unexpected backgrounds. Such was the case in LZ's first science run (SR1), in which several types of novel backgrounds were removed using a wide variety of cuts. This work attempts to use autoencoders (AEs) on PMT waveforms to determine which events are most likely backgrounds.

The LZ Time Projection Chamber (TPC)

LZ searches for dark matter by using PMTs to detect two signals produced by particles scattering on xenon atoms within the TPC:

- S1: scintillation light detected immediately after the scatter
- S2: electroluminescence light produced by electrons after drifting up to the top of the TPC

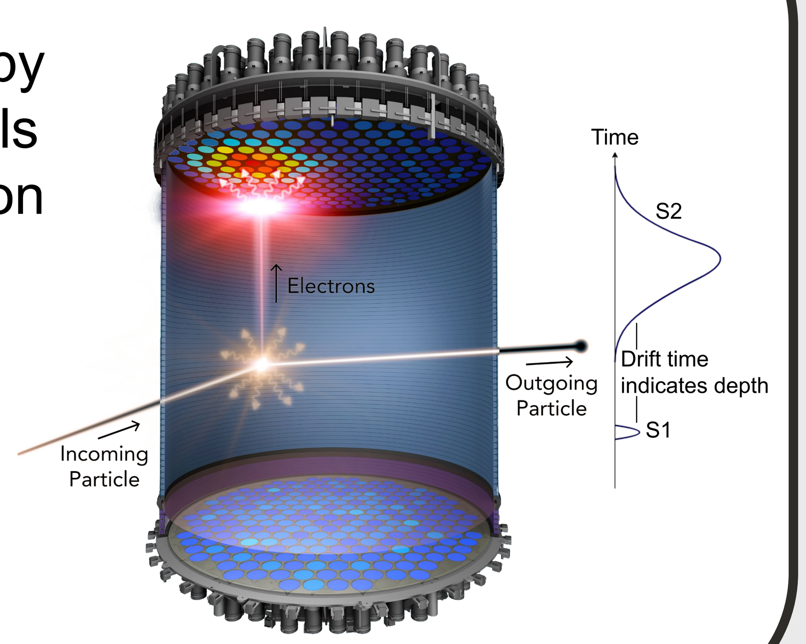

Data Selection and Preparation

This work uses S2 waveforms as their shapes indicate scatter location useful for finding edge backgrounds

Training S2 data:

- SR1 (late 2021 - mid 2022)
- Areas > 600 phd
- Spatially uniform
- Must pass data quality cuts

Waveform preparation:

- Downsampled to 200 MHz
- Normalized by max amplitude
- Smoothed
- Zero-padded to 500 samples

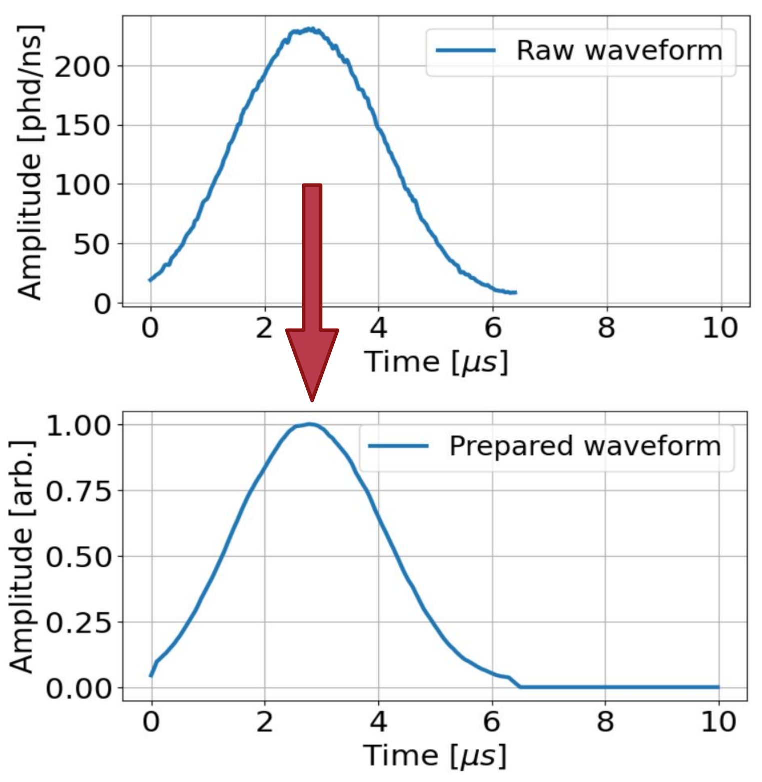

Autoencoder Model

- An autoencoder comprises two neural networks:
  - An encoder: reduces the dimension of input data to a more compact latent space
  - A decoder: uses encoded data to recreate the input data
- Autoencoders for anomaly detection
  - AE performance is worst for inputs which look the least like training inputs - these are "anomalous"
- Tested using either dense layers or 1D convolutional layers
  - Performance was similar, and the latter is used for this work

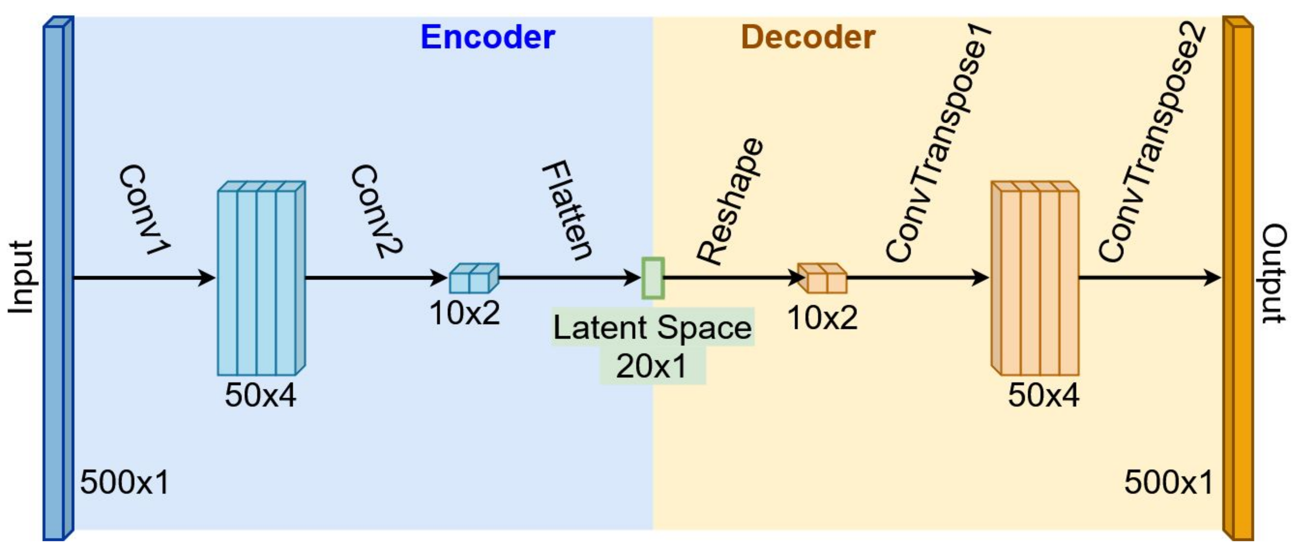

Training

- Mean absolute error (MAE) losses greater than 95% of all quality S2s are labeled anomalous
- By training on quality S2s, the AE struggles to recreate non-quality S2s.

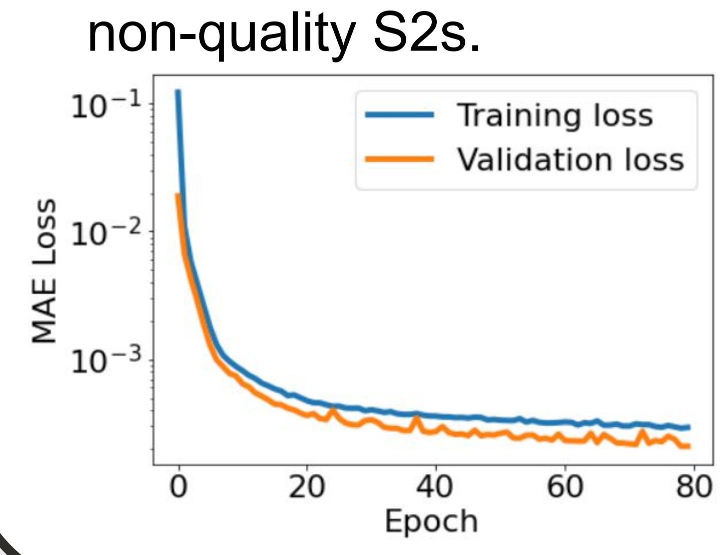

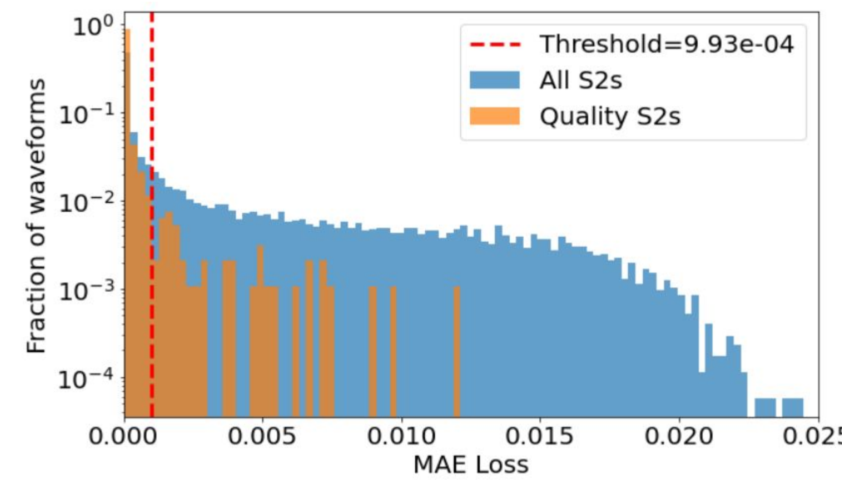

Anomalous Quality S2s

One class of anomalous S2 in the training dataset looks inconsistent with single scatters. This method is able to flag these as likely backgrounds.

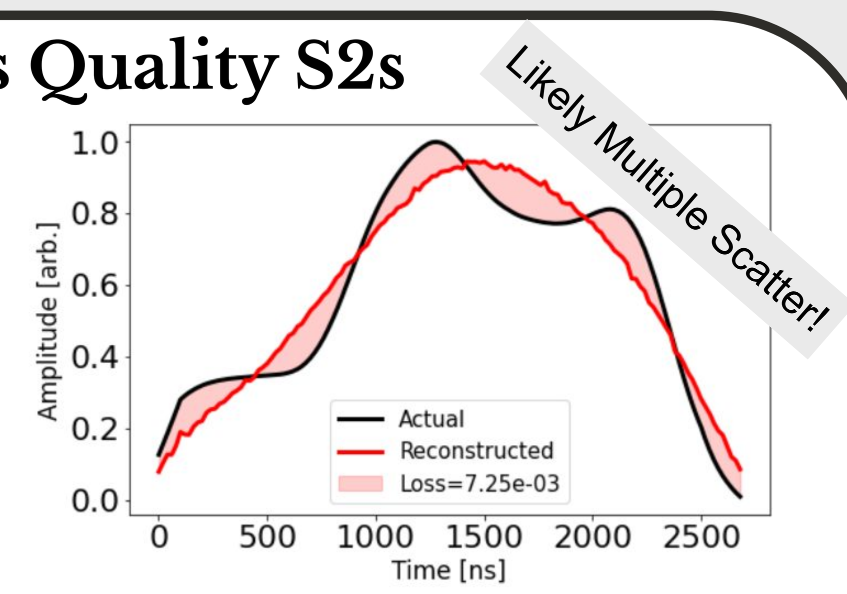

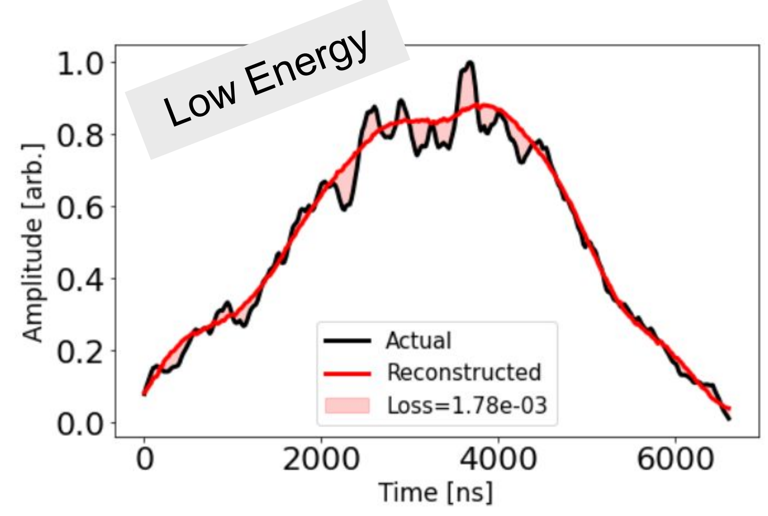

The second class comprises jagged low energy S2s. These are not necessarily backgrounds and show the limitations of this method with low energy S2s.

Performance on all S2 Waveforms

Testing on all single scatters, 42% of S2s are labeled as anomalous (i.e. likely background)

- Nearly all anomalies come at low drift time
- S2s with low area or near the TPC wall tend more anomalous

No S1 information is required here

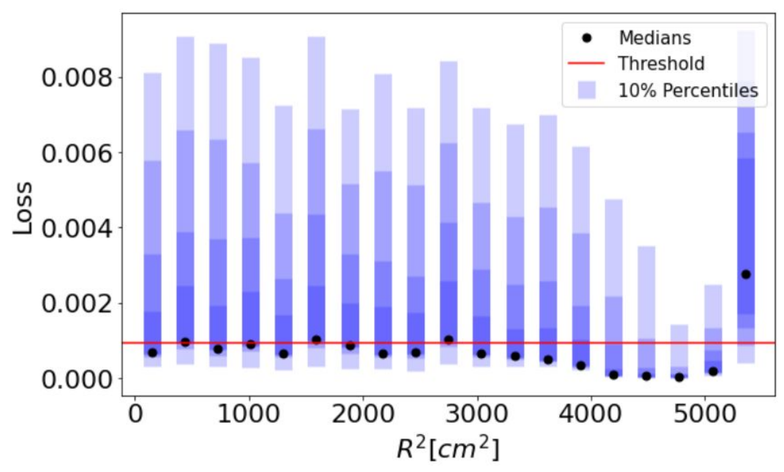

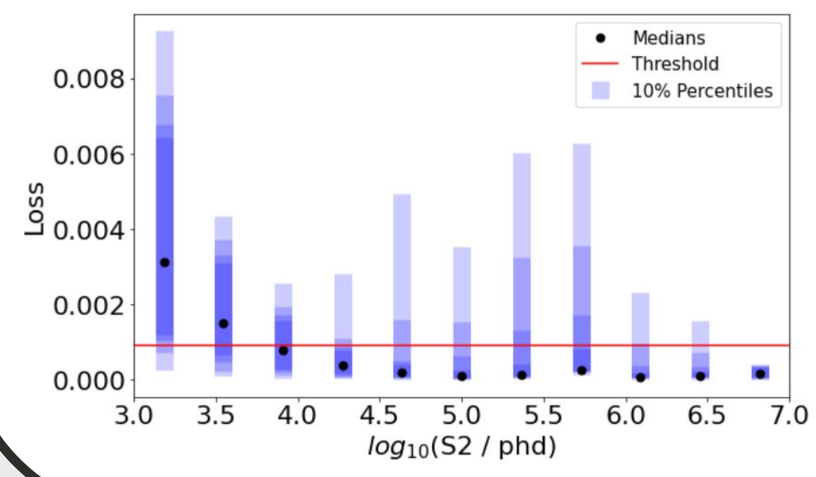

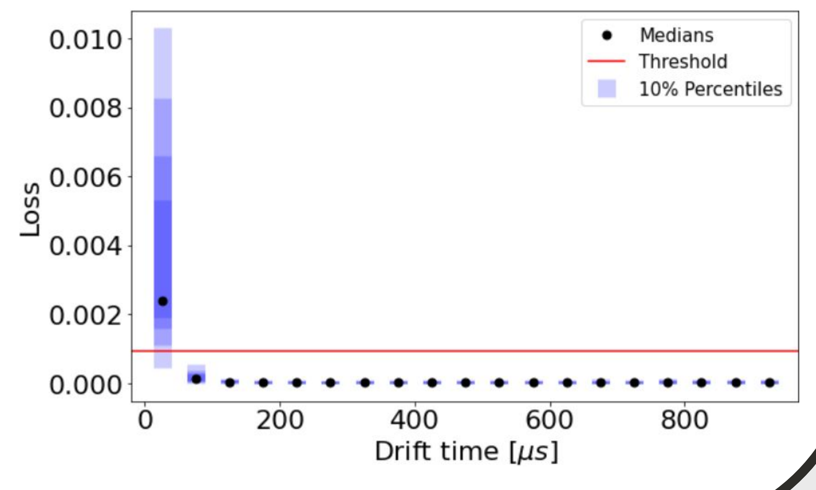

Comparison to SR1 Cuts and Conclusions

While 42% of S2s are marked as anomalous, 56% of S2s are cut by data quality cuts. Overall, the AE labeled 70% of poor-quality S2s as anomalous and 96% of quality S2s as normal.

- Most success came replicating cuts meant to remove scatters near the liquid surface
- Events cut using drift time information were not labeled anomalous as often as the AE does not know the drift time

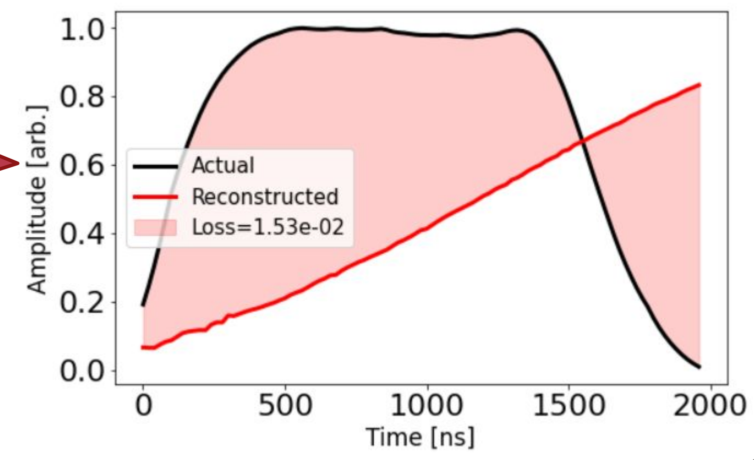

Supported by KIPAC, UKRI Science and Technology Facilities Council, FCT Fundação para a Ciência e a Tecnologia, Sanford Underground Research Facility, IBS Institute for Basic Science