Automating ATLAS control room anomaly detection with ML

Avery Hanna | CERN Summer Student, Tufts University
Mario Campanelli | UCL & Antoine Marzin | CERN

Introduction

To ensure high-quality data acquisition at ATLAS, the detector status is monitored by a team of shifters in the control room. We aim to improve the quality of anomaly detection and decrease the workload of the control room staff by developing a machine learning model to watch the incoming time-series data on the status of the detectors and flag anomalies. The goal is for this model to run online, alerting staff of problems in real-time so the appropriate corrections can be made.

Model architecture

Input:
- Time series data of 5 samples
- Each sample with 16 features: L1 trigger rates, L1 muon sector logic inputs, pileup

Output:
- Prediction of 16 features for next sample (one luminosity block/one minute in future)

Predictive LSTM Autoencoder: Long short-term memory (LSTM) layers incorporate time-series element of data, handling anomalies evolving in time. Autoencoder shape forces the model to learn lower dimensional representation of data, extracting key points from the feature list. Since learning is unsupervised, we can easily retrain on new datasets to adapt to changes in operational conditions.

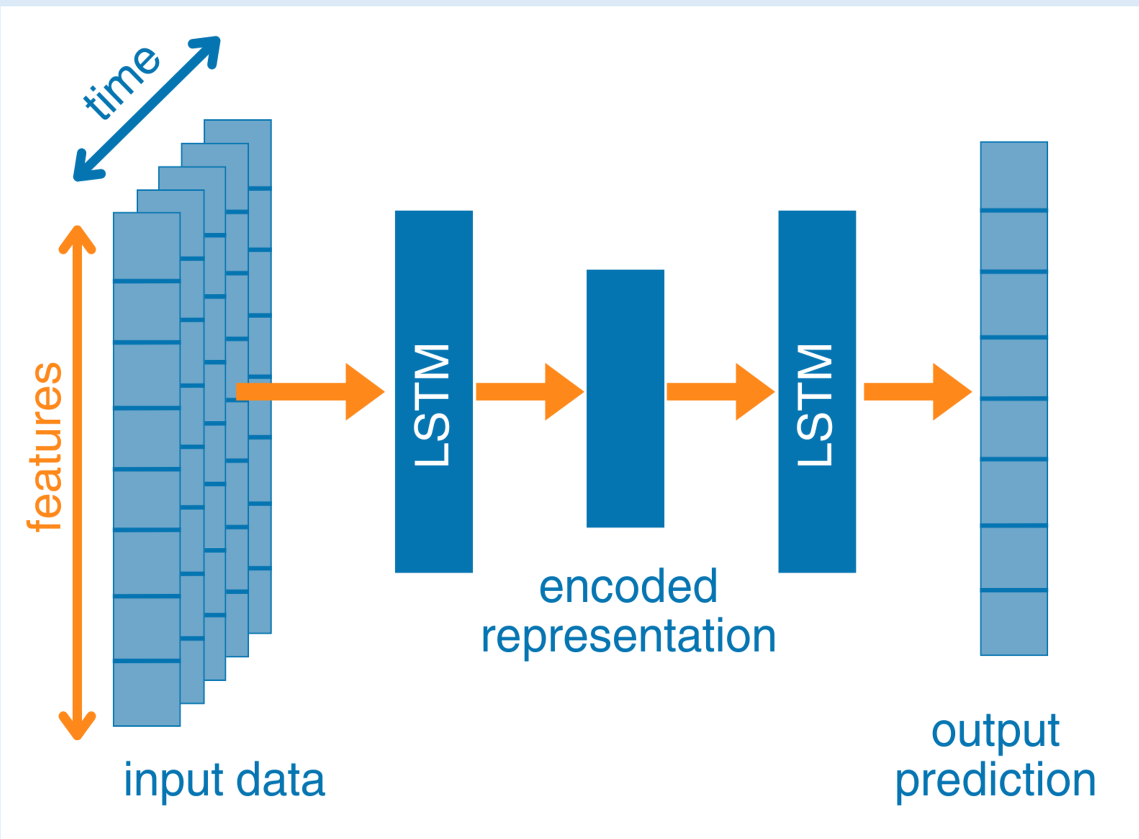

Figure 1. Structure of model with data flowing from input on left to output on right.

Anomaly/clean classification step:
- Model trained exclusively on clean data so low mean squared error (MSE) between predicted and clean real values
- Doesn't see anomalous data in training so when encountered, performs worse (higher MSE)
- Set threshold MSE between clean and anomalous classifications

Data processing

- Import raw data and drop high fluctuating regions at ends of run
- Robust scaling of data so that resilient to outliers
- Note luminosity block marking transition from constant to decaying pileup (e.g. LB 700 in Fig. 3) so that data can be separated

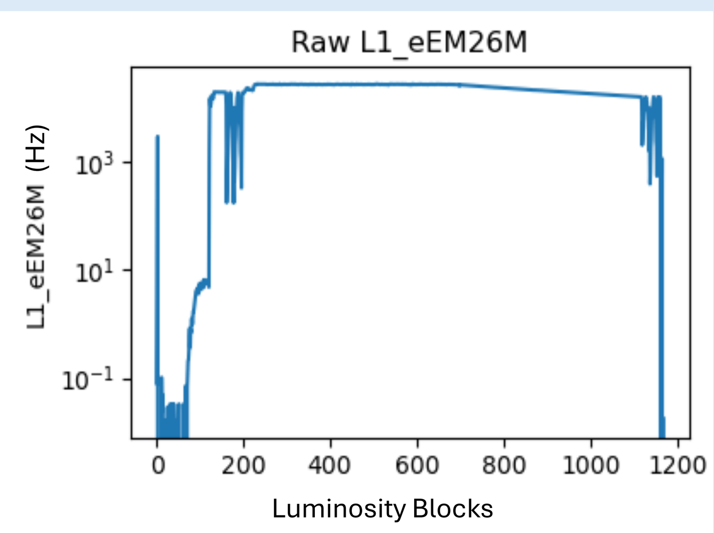

Figure 2. Raw data from run 480219 pre-trimming and scaling.

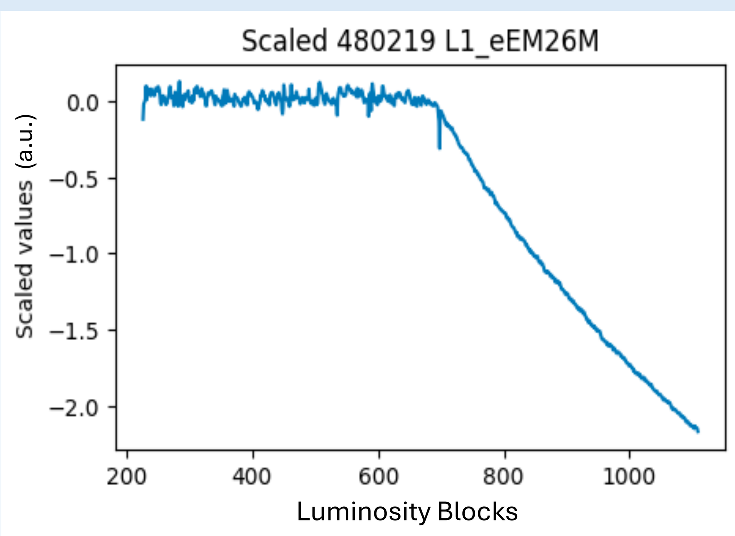

Figure 3. Post-processed data that will be fed into model.

Model performance

Training data:
- Three full runs
- Around 2000 datapoints

Test data:
- Fourth full run
- Around 900 datapoints

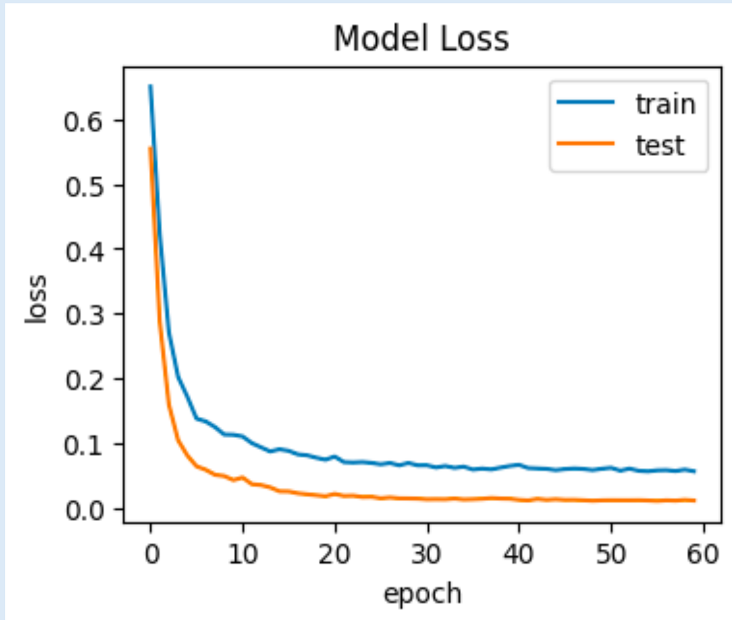

Figure 4. Loss curve shows strong improvement of MSE over epochs.

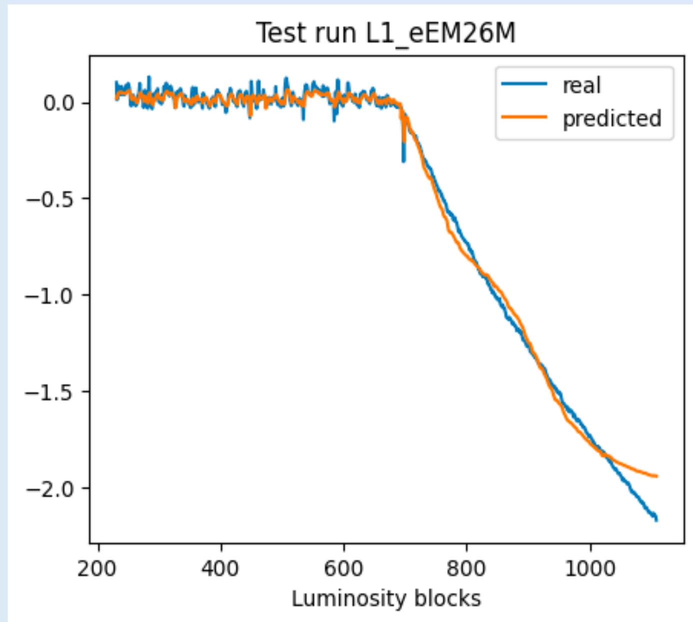

Figure 5. Real test data vs the predictions made by the model.

Model training
- Loss = MSE
- Test sample consistently performing better than training results from intrinsic differences in test and train

Model predictions
- Model accurately traces trends in data
- Prediction smoother than real as intended with autoencoder
- High MSE at ends of runs because fails to reach lows in low pileup region since very few samples in that region

Detecting 5% anomaly

To assess the model's effectiveness at identifying anomalies, we produced artificial anomalous data sets by taking the test run and, for a single feature, increasing the first 30 data points by 5%.

- Most of the high MSE data points belong to anomalous data range but also observe bump at end of data where prediction fails to keep up with drop in data
- Can set threshold to get 90% true positive rate at cost of around 30% false positive rate
- Some features have poor ROC curves, but at 10% anomaly, AUC greater than 0.66 for all datasets

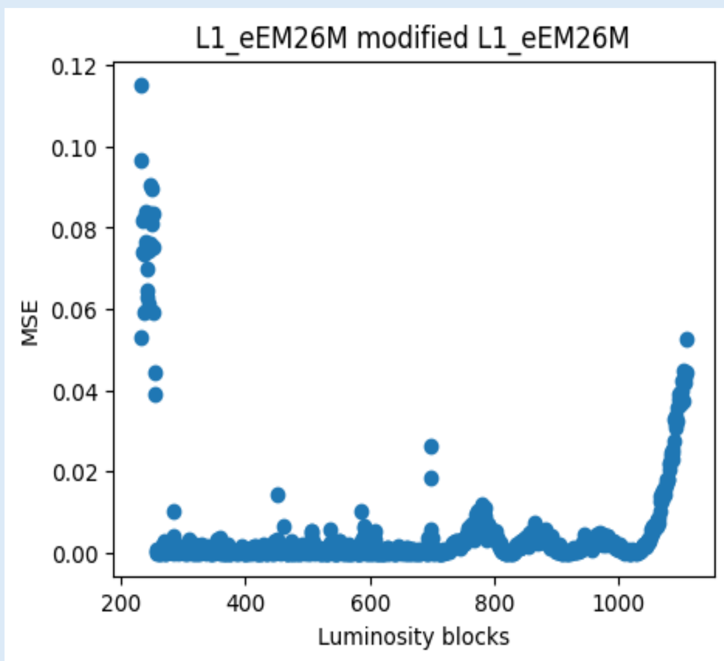

Figure 6. MSE between predicted and real values of L1_eEM26M for test data with this feature modified for first 30 datapoints.

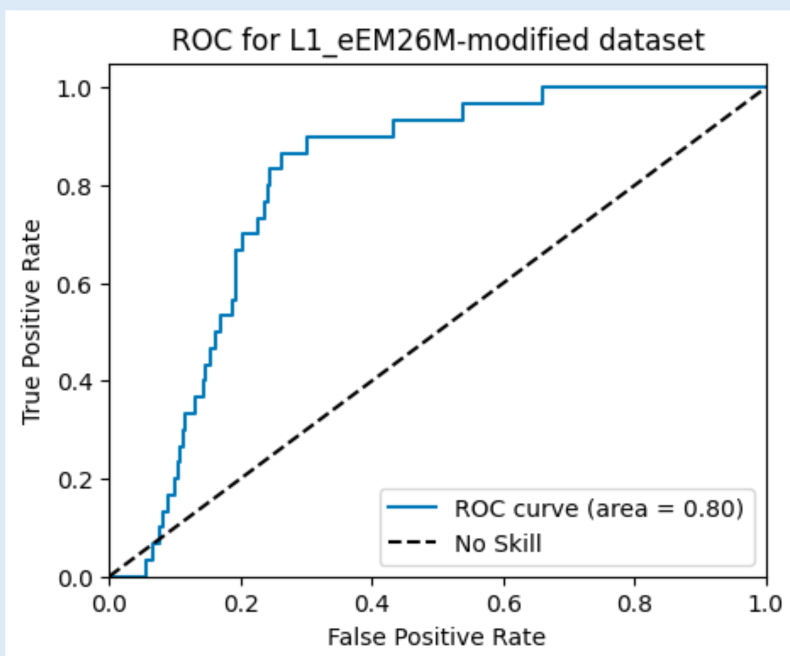

Figure 7. ROC curve shows trade-off between true positive rate and false positive rate as classification threshold is adjusted.

Next steps

- Assess different classes of anomalies (e.g. continuous anomalous range, point anomalies, varying anomalies)
- Test on real anomalous runs
  - Portion of data from June with muon endcap disabled
- Experiment with pinpointing source of error
- Explore possibilities of variational autoencoder architecture
- Prepare for online usage

avery.bryn.hanna@cern.ch