ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL013

INITIAL RESULTS OF APPLYING AN AUTOENCODER TO DETECT ANOMALIES IN THE AIR CONDITIONING SYSTEMS OF THE BROOKHAVEN ACCELERATOR COMPLEX*

Y. Gao†, I. Blackler, K. Brown[1], J. Morris, V. Schoefer, R. Terheide,
C-AD Dept., Brookhaven National Laboratory, Upton, NY, USA.
B. Huang, AMS Dept., Stony Brook University, Stony Brook, NY, USA.
[1]also at ECE Dept., Stony Brook University, Stony Brook, NY, USA

Abstract

The Brookhaven National Laboratory (BNL) Collider-Accelerator Complex contains millions of control points. Monitoring tolerances for these control points is crucial for the system and is a challenging task. Catching early signs of failures in those systems will be very beneficial as they can save extensive downtime. Anomaly detection in particle accelerators has been highlighted and can significantly impact system performance. Autoencoder is one of the most commonly used techniques for detecting anomalies. In this contribution, we apply an autoencoder method to analyze the historical data for runs 21 and 22 to find precursors for trips (and actual trips) of Air Conditioning (AC) systems based on local thermostat readbacks. Results from the existing system are presented, showing that the new method can catch early signs of AC trips so that advance notices can be sent for the operators to take prompt action.

INTRODUCTION

The Brookhaven National Laboratory (BNL) Collider-Accelerator Complex spans over two square miles and comprises thousands of elements. The Air Conditioning (AC) system is part of the critical infrastructure for the Complex. There are many different AC units with widely varying performance characteristics. Many of these AC units are not network connected and hence cannot report their status remotely. The only readbacks available remotely are from the thermostats used to control the units and other thermometers scattered throughout the Complex. The heat loads in these buildings can vary dramatically, depending on the power loads on the various pieces of equipment stored within them, so wide temperature swings can be expected.

Currently, alarms are generated for high and low-temperature readings for these locations, but they are set with large tolerances to prevent false alarms. Consequently, an AC unit can trip, and the temperature rises steadily for multiple hours before an alarm is generated. The drawback of this approach is that it misses the opportunity to catch anomalies early on and depends too much on the threshold value. Catching trips of the AC units promptly–or even in advance–allows for more timely fault resolution. It also reduces the risk of equipment damage and extended downtime.

In this work, we apply an autoencoder approach. The purpose is to develop a more intelligent way to catch precursors of anomalies ahead of time, so that operators can take proactive actions. Initial results of anomaly detection on some AC historical data are presented.

Autoencoder [1] is a common technique to detect anomalies from the input data. It is capable of learning a dense representation of the data, called latent representations or codings, in an unsupervised way (i.e. the training set is unlabeled). The latent space usually has a much lower dimensionality than the original data, hence an autoencoder is also good for dimensionality reduction purposes.

A typical autoencoder structure is shown in Fig. 1. It is symmetric and contains an encoder and decoder. The encoder abstracts the input data into a smaller latent space, which contains the high-level information of the data. The decoder decodes the latent space into output and tries to recover the input. For anomaly detection, reconstruction errors between every pair of input and output are computed, and an error threshold is set. When new data is available, new reconstruction errors will be computed. An error higher than the predefined threshold signals a new data pattern. Depending on the application, it could be a new data logic or an anomaly. By controlling the length of the input sequence, different data patterns can be captured.

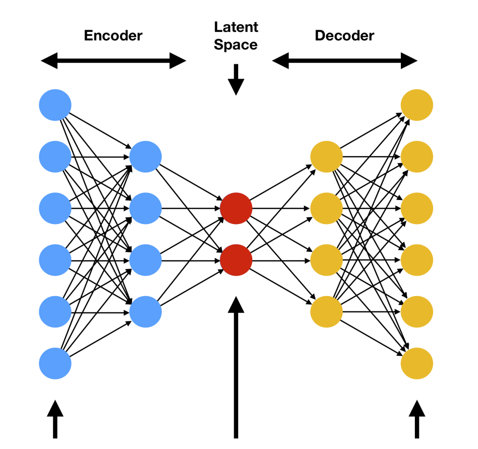

Figure 1: A typical structure of an autoencoder.

Anomaly detection techniques have been widely applied in the accelerator fields to improve operations. Work [2–4] propose schemes to understand and predict faulty behavior

* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.
† ygao@bnl.gov

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL013

in superconducting RF cavities and magnets [5]. Work [6] uses Machine Learning (ML) to identify and remove malfunctioning beam position monitors in the LHC. Work [7] applies ML to detect errors in hardware installation.

EXPERIMENTAL RESULTS

In this work, we test the autoencoder on historical building temperature data. The original datasets are taken from support buildings circling the Alternate Gradient Synchrotron (AGS) ring. The buildings are old, as are most of their AC systems. Hence they are more likely to exhibit anomalous behavior.

Since we are only interested in capturing the data pattern (not the individual data values), for computing efficiency reasons, the raw data is downsampled by a factor of 10. Then they are standardized and scaled in favor of the neural network so that the datasets have a 0 mean and unit standard deviation (std). Moreover, to exclude the influences from the daily weather fluctuations, every dataset is regulated by subtracting the average temperatures of 24-hour windows. The datasets after those preprocessing steps are shown in Fig. 2.

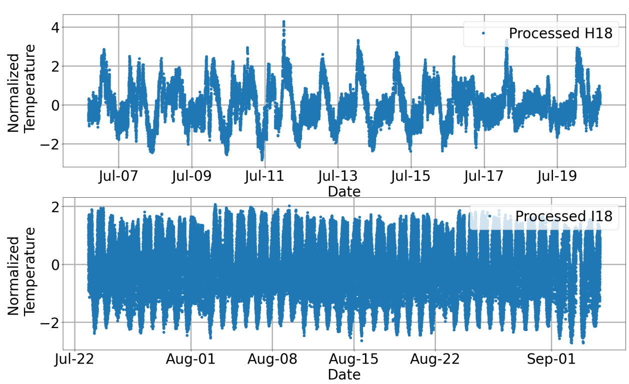

Figure 2: Processed building H18 and I18 temperature data. The data are downsampled, normalized, and shifted to compensate for the influences from the ambient environment.

Next, we identify the input sequences to train the autoencoder. In other words, this is the pattern the autoencoder will learn and based on which to predict anomalies. The prediction process works as follows. On the left of Fig. 3, it shows the pattern the autoencoder trained to capture. It works like a shifting window. New data comes in, and old data moves out. When the first different point comes to join the input sequence, the whole data pattern changes, and an anomaly is detected. The first different point could be a precursor of an anomaly (or not). That's better than the traditional hard-threshold method, where the data could already go wrong before it reaches the predefined value. For either building dataset, we use a 30-minute long sequence that can cover an entire period of a temperature waveform as the input.

The autoencoder has a symmetric network topology with 32, 16 nodes in middle layers, 10 nodes for latent space, and variable encoded dimension depending on the input data.

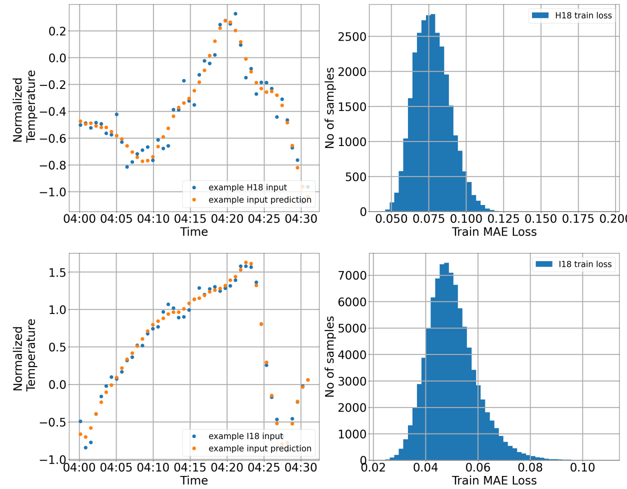

Figure 3: Autoencoder input sequences are shown on the left for each dataset (blue points). A 30-minute long time sequence is selected, which can cover an entire period of a temperature waveform for either building. It serves as an input instance for the autoencoder. After training, the autoencoder can accurately capture the data pattern (orange points). The Gaussian-shape distribution of training errors on the right also validates the performance.

ELU [8] is adopted as the activation function with He-normal kernel initialization [9]. After training, the autoencoder can recover the input patterns pretty well, as shown in Fig. 3 (left). The orange points are the neural network predictions, which match the training data (blue) with good accuracy.

We use the Mean Absolute Error (MAE) as the reconstruction error [10]. The distribution of the train MAE loss is shown in Fig. 3 (right). The Gaussian-shape distributions of training errors indicate that the autoencoder performs well on the training set.

The test data for each building are picked from a different time period, as shown in Fig. 4. We can see some anomalous patterns, such as the large high-temperature waves in Building H18 and high spikes in Building I18.

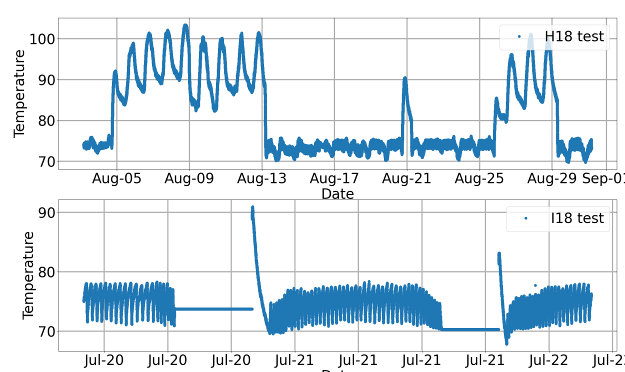

Figure 4: The test data for each building are picked from a different time period.

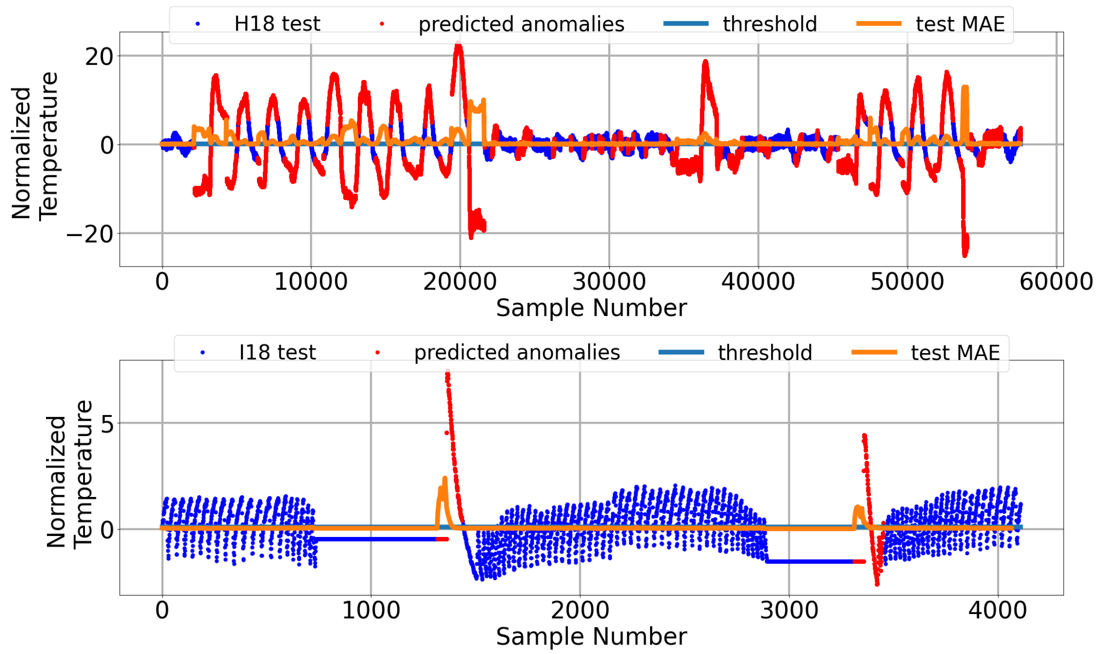

Figure 5: Autoencoder results on the test data. It successfully detects data anomalies. However, more tuning is needed to improve the accuracy around the borders between different groups, and to avoid overfitting.

The test data are processed using the same steps as the training data and are fed to the autoencoder to detect anomalies. The results are shown in Fig. 5. The blue plots are the processed test data, and the predicted anomalies are marked in red. The threshold is set to be the maximum training reconstruction error, and the test reconstruction errors are shown in orange. Whenever the test errors pass the threshold, the corresponding data sequence will be marked as anomalies.

We can see that the autoencoder has successfully detected the anomalies on those large spikes, which matches our expectations. However, there are several improvements to be made. First, the shifting introduces discontinuities in the data waveform, which leads to false-positive predictions around the border areas between different data groups. Second, due to overfitting [11], the model ignores anomalies with simpler patterns (e.g., the straight lines on July 20 and 21) where it can reconstruct well.

FUTURE WORK

First, we would like to tune the algorithm to improve its accuracy and eliminate the wrong predictions.

Second, we would like to try the dataset with a variational autoencoder [12]. The difference between an autoencoder and a variational autoencoder is that the traditional autoencoder generates a latent vector while a variational autoencoder learns to generate two vectors that represent the parameters (mean and variance) of a distribution from which the latent vector is sampled, and which the decoder can transform back to the original input. Therefore, its training is more regularised to avoid overfitting.

Third, after we assess the performance of different algorithms, we would like to implement the final algorithm in the actual system and try with live data from operations.

CONCLUSION

In this work, we apply an autoencoder to help detect and predict anomalies from building temperature data. The results show that the autoencoder has successfully detected anomalies in the test data.

ACKNOWLEDGEMENTS

The author wants to thank Ian Blackler for providing real machine data for experiments. The author wants to thank Rachel Terheide and Bohong Huang for helping analyze the data.

REFERENCES

[1] J. Schmidhuber, "Deep learning in neural networks: An overview," *Neural Networks*, vol. 61, pp. 85–117, 2015.

[2] A. S. Nawaz, S. Pfeiffer, G. Lichtenberg, and H. Schlarb, "Self-organzied critical control for the european xfel using black box parameter identification for the quench detection system," in *2016 3rd Conference on Control and Fault-Tolerant Systems (SysTol)*, 2016, pp. 196–201. doi:10.1109/SYSTOL.2016.7739750

[3] A. Nawaz, S. Pfeiffer, G. Lichtenberg, and P. Rostalski, "Anomaly detection for the european XFEL using a nonlinear parity space method," *IFAC-PapersOnLine*, vol. 51, no. 24, pp. 1379–1386, 2018, 10th IFAC Symposium on Fault Detection, Supervision and Safety for Technical Processes SAFEPROCESS 2018. doi:10.1016/j.ifacol.2018.09.554

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL013

[4] A. Nawaz, G. Lichtenberg, S. Pfeiffer, and P. Rostalski, “Anomaly Detection for Cavity Signals - Results from the European XFEL,” in *Proc. 9th International Particle Accelerator Conference (IPAC’18), Vancouver, BC, Canada, April 29-May 4, 2018*, Vancouver, BC, Canada, 2018, pp. 2502–2504. doi:10.18429/JACoW-IPAC2018-WEPMF058

[5] M. Wielgosz, A. Skoczeń, and M. Mertik, “Using LSTM recurrent neural networks for monitoring the lhc superconducting magnets,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 867, pp. 40–50, 2017. doi:10.1016/j.nima.2017.06.020

[6] E. Fol, J. C. de Portugal, and R. Tomás, “Unsupervised Machine Learning for Detection of Faulty Beam Position Monitors,” in *Proc. 10th International Particle Accelerator Conference (IPAC’19), Melbourne, Australia, 19-24 May 2019*, Melbourne, Australia, 2019, pp. 2668–2671. doi:10.18429/JACoW-IPAC2019-WEPGW081

[7] T. Dewitte, W. Meert, P. V. Trappen, and E. V. Wolputte, “Anomaly Detection for CERN Beam Transfer Installations Using Machine Learning,” in *Proc. ICALEPCS’19*, New York, NY, USA, 2020, pp. 1066–1070. doi:10.18429/JACoW-ICALEPCS2019-WEMPR010

[8] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, *Fast and accurate deep network learning by exponential linear units (ELUs)*, 2015. doi:10.48550/arXiv.1511.07289

[9] K. He, X. Zhang, S. Ren, and J. Sun, *Delving deep into rectifiers: Surpassing human-level performance on imagenet classification*, 2015. doi:10.48550/arXiv.1502.01852

[10] *Mean absolute error*. https://en.wikipedia.org/wiki/Mean_absolute_error

[11] *Overfitting*. https://en.wikipedia.org/wiki/Overfitting

[12] D. P. Kingma and M. Welling, *An introduction to variational autoencoders*, 2019. doi:10.48550/arXiv.1906.02691