ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-WEPA104

PREDICTION OF SUPERCONDUCTING MAGNET QUENCHES WITH MACHINE LEARNING

M. Kilpatrick*, J. Edelen[1], J. Einstein-Curtis[1], R. O'Rourke[1], RadiaSoft LLC, Boulder, USA
K. Drees[2], M. Valette[2], Brookhaven National Laboratory, Upton, USA

Abstract

Superconducting magnet technology is one of the foundations of large particle accelerator facilities. A challenge with operating these systems is the possibility for the magnets to quench. The ability to predict quenches and take precautionary action in advance would reduce the likelihood of a catastrophic failure and increase the lifetime operability of particle accelerators. We are developing a machine learning workflow for prediction and detection of superconducting magnet quenches. In collaboration with Brookhaven National Laboratory (BNL), our methods for algorithm development will utilize magnet data from test stands and the Relativistic Heavy Ion Collider (RHIC) ring magnets to allow for a robust identification of magnet quenches. Our methods divide the problem into two different aspects. First, we are developing machine learning algorithms for binary and multi-classification of the various types of quench events. Second, our prototype machine learning model will be used to predict a quench event using precursor identification. We plan to integrate and test our monitoring system at the BNL facility to perform quench identification and prediction.

BACKGROUND

Quench protection systems [1–4] are used to prevent potentially catastrophic failures in superconducting magnet systems. There is an extensive protection system in place at BNL for the superconducting magnets in RHIC. The conventional systems have two primary mechanisms, actively monitoring the resistance of the superconducting cable and providing a relief system for the cryogenic system. The conventional quench protection system constantly monitors the voltages and currents of the power supplies (PS) on the ring to insure the resistances are below a threshold. When a quench occurs the quench protection system ensures the magnets' safety while the beam abort system extracts the beam to prevent damage to the accelerator. Due to the latency between these systems however, the beam extraction from a quench can still result in potential damage to the accelerator. Our aim is to detect the onset of a quench before it happens and trigger the abort accordingly to avoid further issues. We are developing our algorithms on quench data from magnets in the RHIC.

DATA PREPARATION

We have acquired upwards of a decade of abort data from operational magnets at RHIC. Each of these years has numerous triggered events which could be a quench or other possible aborts caused by RF, beam loss, etc. When an abort occurs the data logger records a snapshot of the machine state leading up to and immediately following the abort. From each of these readouts, we have PS and beam position monitoring (BPM) data from the magnets in the yellow and blue rings of the accelerator. Once a quench event is triggered, the beam position, difference in position, and coherence are readout from the BPM at a rate of 10 kHz. At the same time, the reference currents, measured currents, voltages, and error of the voltage are readout from the PS at a rate of 720 Hz. It should be noted that the PS of the rings are connected to varying magnets along the ring of the accelerator which causes an inherent uncertainty in the location of the quench.

Our current methods combine all of these files into a HDF5 [5] format, such that the data is easily parsed, contains metadata, and can have a multitude of different selections for a variety of parameter training and classification schemes. We have developed a python interface to read the HDF5 data files and return useful formats for input into machine learning pipelines.

CLASSIFICATION

First we developed classifiers to label new data based historical training data. Here we consider binary classifiers that label if a quench occurred or not and classifiers that can label the type of abort. Additionally we consider classifiers that analyze data for the entire ring and device level classifiers that look at single magnets or sequences of magnets. With our data parser described above, we have a multi-classification of the datasets which is used as a label for input into a boosted decision tree (BDT) [6] and a custom multilevel perceptron (MLP) [7].

In figures 1 and 2, we see that both methods are able to distinguish a quench from all other labels using the entire time sequence of PS data. The BDT method, which defines a set of nodes and leaves to split data with the goal of maximizing the information gained from the data or minimizing entropy shows the best area under the curve of 0.88.

UNSUPERVISED LEARNING

We have developed vanilla autoencoders (AE), see fig. 3, and a Long Short-term Memory (LSTM) [8] autoencoders for the identification of quench precursors within the PS datasets. For the purposes of this proceeding, we will concentrate on the LSTM autoencoder which is trained and validated with PS data from RHIC. To identify quench precursors and the dynamic structure of a quench event the PS data is split into a pre-quench and quench category. Each PS dataset contains data for three seconds leading up to the

* kilpatrick@radiasoft.net

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-WEPA104

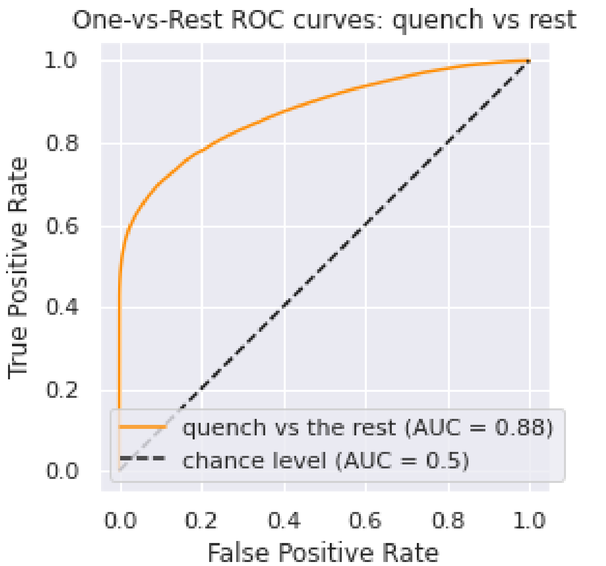

Figure 1: The ROC curve for the quench vs. rest using a BDT with PS data from BNL.

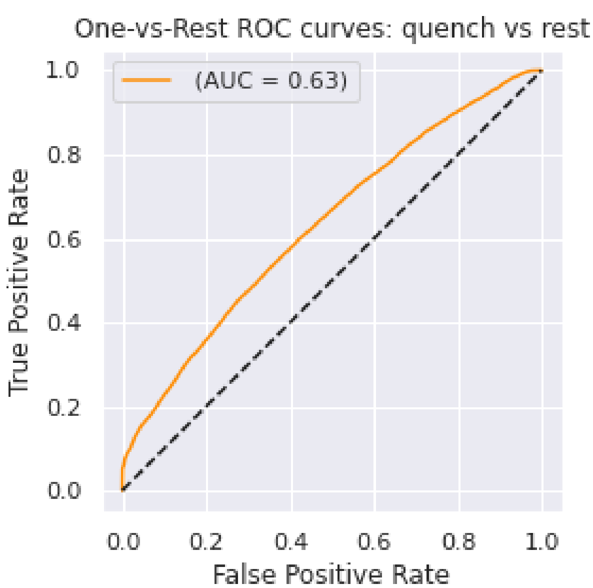

Figure 2: The ROC curve for the quench vs. rest using a MLP with PS data from BNL.

abort and one second afterwards. The abort event occurs at approximately time step 2172 in each dataset.

For training, the data were split in time using a 2:1 ratio while removing the last 100 ms leading up to the abort trigger. The autoencoder is trained and validated on the first two-thirds of the data and tested on the last third. The 100ms buffer allows time for our precursor identification tool to trigger an abort before the quench occurs. Once the data were split into training, validation, test, we further segmented the data by sequence length. This sets the number of clock cycles that will be seen by the autoencoder at a given time and is referred to as the precursor window. Initial scans were performed to optimize the precursor window, which

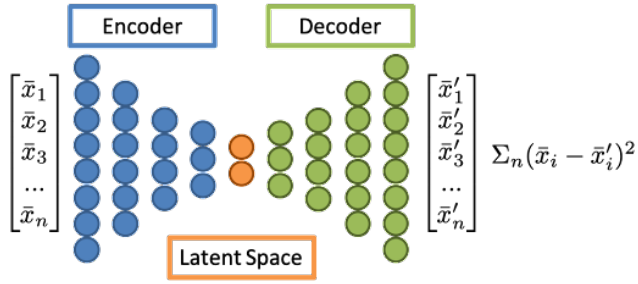

Figure 3: Architecture for an autoencoder which shows the decrease in dimensionality to a latent space along with a decoding architecture back to the original input size.

trained multiple models on a range of data slices from 14 to 700 clock cycles. A metric to maximize the difference in the reconstruction accuracy of pre-quench and quench data determined the optimal window size to be 660 clock cycles for PS data.

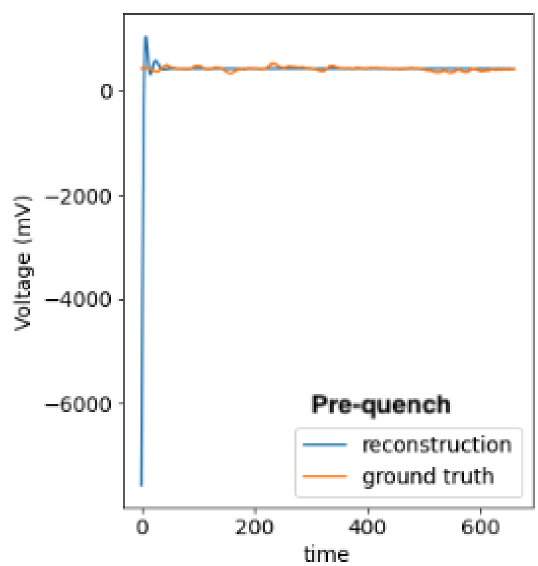

Figure 4: The reconstruction and ground truth of a pre-quench slice of a PS that was classified as a quench.

Now we have an optimized precursor window that should have a maximized difference in reconstruction accuracy for pre-quench and quench events. In figures 4 and 5, the reconstruction of a magnet pre-quench and magnet quench show the large difference in accuracy.

Examining the LSTM latent space, see fig. 6 and 7, we see a transform of the 10 dimensional latent space to a 2D representation with TSNE [9] and principle component analysis [10], respectively. These dimensionality reduction techniques show a promising methods to identify quench events for future runs within RHIC. The latent space representations show a separation of pre-quench, quench, and no-quench (a validation dataset). We are investigating methods of clustering on these representations to accurately identify features of quench precursors.

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-WEPA104

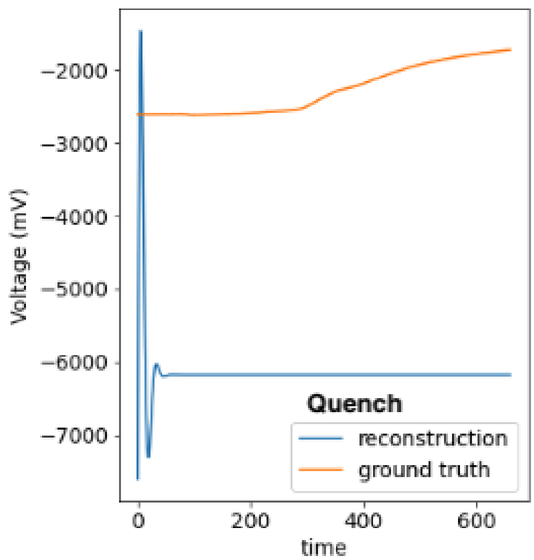

Figure 5: The reconstruction and ground truth of a quench slice of a PS that was classified as a quench.

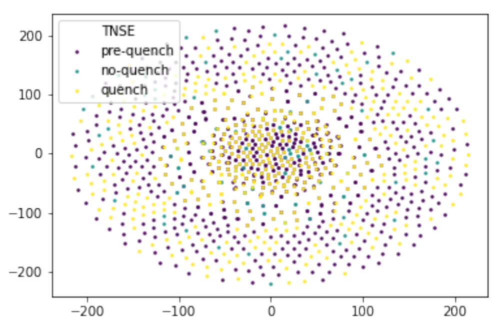

Figure 6: The LSTM latent space after transforming to a 2D manifold with TSNE.

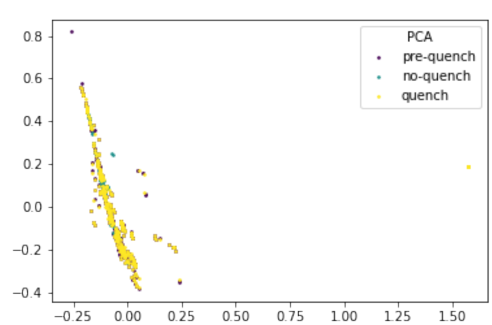

Figure 7: The LSTM latent space after transforming to 2 components with a PCA.

SUMMARY

We have developed a machine learning framework for quench classification and precursor identification. Our results on classification show that there are distinguishing features for PS data that allow for automatic methods for operators to classify events. We demonstrated that a LSTM autoencoder can be trained on pre-quench PS data and can identify quench precursors due to predicted reconstruction inaccuracies. From these models, the latent space shows potential for clustering algorithms to identify regions for a further ability to distinguish quenches from non-quenches.

ACKNOWLEDGEMENTS

This research used resources at the Relativistic Heavy Ion Collider, a DOE Office of Science User Facility operated by Brookhaven National Laboratory. This material is based upon work supported by the U.S. Department of Energy, Office of Science. Phase I SBIR under Award Number(s) DE-SC0022795.

REFERENCES

[1] P. F. Smith. “Protection of Superconducting Coils”. In: Review of Scientific Instruments 34.4 (1963), pp. 368–373. eprint: https://doi.org/10.1063/1.1718368.

[2] C. Conkling. “RHIC beam permit and quench detection communications system”. In: Proceedings of the 1997 Particle Accelerator Conference (Cat. No.97CH36167). Vol. 2. 1997, 2496–2498 vol.2.

[3] L. Salasoo. “Superconducting magnet quench protection analysis and design”. In: IEEE Transactions on Magnetics 27.2 (1991), pp. 1908–1911.

[4] J. P. Edelen and N. M. Cook. “Anomaly Detection in Particle Accelerators using Autoencoders”. In: arXiv preprint arXiv:2112.07793 (2021).

[5] The HDF Group Hierarchical Data Format, version 5. (https://www.hdfgroup.org/HDF5/)

[6] J. Friedman, Greedy function approximation: a gradient boosting machine. *Annals Of Statistics*. pp. 1189-1232 (2001)

[7] Haykin, S. Neural networks: a comprehensive foundation. (Prentice Hall PTR,1994)

[8] Y. Ma, et al. Long Short-Term Memory Autoencoder Neural Networks Based DC Pulsed Load Monitoring Using Short-Time Fourier Transform Feature Extraction. *2020 IEEE 29th International Symposium On Industrial Electronics (ISIE)*. pp. 912-917 (2020,6), ISSN: 2163-5145

[9] L. Maaten, et al. Visualizing Data using t-SNE. *Journal Of Machine Learning Research*. **9**, 2579-2605 (2008), http://jmlr.org/papers/v9/vandermaaten08a.html

[10] I. Jolliffe, et al. Principal component analysis: a review and recent developments. *Philosophical Transactions Of The Royal Society A: Mathematical, Physical And Engineering Sciences*. **374**, 20150202 (2016,4), https://royalsocietypublishing.org/doi/10.1098/rsta.2015.0202, Publisher: Royal Society