ISBN: 978-3-95450-232-5 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2022-WEPA29

REAL-TIME CAVITY FAULT PREDICTION IN CEBAF USING DEEP LEARNING*

Md. M. Rahman†, K. Iftekharuddin
Vision Lab, Department of Electrical and Computer Engineering,
Old Dominion University, Norfolk, VA 23529, USA
A. Carpenter, T. McGuckin, C. Tennant, L. Vidyaratne
Thomas Jefferson National Accelerator Facility, Newport News, VA 23606, USA

Abstract

Data-driven prediction of future faults is a major research area for many industrial applications. In this work, we present a new procedure of real-time fault prediction for superconducting radio-frequency (SRF) cavities at the Continuous Electron Beam Accelerator Facility (CEBAF) using deep learning. CEBAF has been afflicted by frequent downtime caused by SRF cavity faults. We perform fault prediction using pre-fault RF signals from C100-type cryomodules. Using the pre-fault signal information, the new algorithm predicts the type of cavity fault before the actual onset. The early prediction may enable potential mitigation strategies to prevent the fault. In our work, we apply a two-step fault prediction pipeline. In the first step, a model distinguishes between faulty and normal signals using a U-Net deep learning architecture. In the second step of the network, signals flagged as faulty by the first model are classified into one of seven fault types based on learned signatures in the data. Initial results show that our model can successfully predict most fault types 200 ms before onset. Our fault prediction model shows poor model performance on fast-developing fault types.

INTRODUCTION

The Continuous Electron Beam Accelerator Facility (CEBAF) at Jefferson Lab is a high power, continuous wave recirculating linear accelerator (linac) servicing four different experimental nuclear physics end stations [1]. CEBAF completed an energy upgrade from 6 GeV to 12 GeV in 2017 which required the installation of 11 additional cryomodules, called C100s for their capability to provide 100 MV of energy gain [2]. A schematic of CEBAF with locations of C100 cryomodules is shown in Fig 1. Each cryomodule is composed of 8 superconducting radio-frequency (SRF) cavities. In addition, a digital low-level radio frequency system (LLRF) is developed to regulate the new cryomodules.

CEBAF experiences frequent short machine downtime trips caused by numerous SRF system faults, especially when cavity gradients are being pushed to their limits. In 2019, CEBAF experienced an average of 4.1 RF downtime trips per hour, culminating in approximately 1 hr of beam time lost each day [3]. A data acquisition system is implemented in the C100 cryomodules to record data to investigate the nature and the origin of the SRF faults. These recorded waveform data are analyzed by a subject matter expert (SME) to determine the cavity that caused the trip and the type of fault. This is a non-trivial, laborious task. Typically, a SME performs this task days or weeks after the events. Previous work successfully addressed this fault classification task with machine learning (ML) [4]. In this work, our goal is to develop deep learning-based artificial intelligence (AI) techniques to predict the fault before its onset.

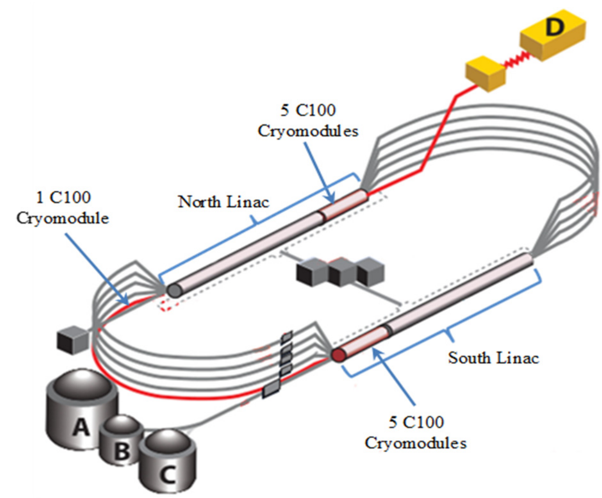

Figure 1: A schematic of CEBAF with the experimental halls (A, B, C, and D), and the locations of the C100 cryomodules labeled [3].

DATA

The data acquisition system in CEBAF synchronously acquires timestamps and saves waveform records of 17 different RF signals from the C100 cryomodule. The data acquisition system includes two primary components, the low-level RF (LLRF) controls and the experimental physics and industrial control system (EPICS), along with a collection of high-level applications. These two components work together to generate and save data for further analysis. Each of the recorded 17 signals is 8192-time steps long. The duration of the recorded signals is approximately 1637.4 ms at a sample interval of 0.2 ms.

There are two types of datasets used for this study. The first type is the normal running dataset which is representative of normal operating conditions (i.e. no faults). For this experiment, we use 60,000 normal running examples of a 100 ms time window. The second type of the

* This work is supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under Contract No. DE-AC05-06OR23177.
† mrahm006@odu.edu

ISBN: 978-3-95450-232-5 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2022-WEPA29

dataset is the faulty data collected from RF cavity trip events which is presented in Fig. 2. Note that approximately 94% of the captured waveform ($t < 0$) represents pre-fault activity, with $t = 0$ the fault onset, and $t > 0$ the post-fault data. Table 1 summarizes the dataset composition with respect to fault types. We use 4,983 faulty events for this analysis. In this experiment, we use just 4 of the 17 recorded signals (GMES, GASK, CRFP, DETA2), identified by SMEs as having the greatest predictive power. We perform min-max normalization for each channel of the signal which transform signal values between 0 and 1. We utilize a 100 ms window to predict the impending faults.

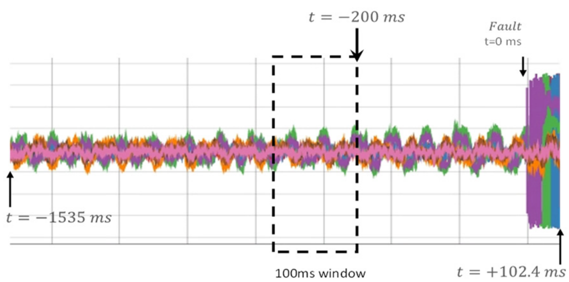

Figure 2: Waveform captured by the data acquisition system during a fault event. The total duration of the waveform is 1.64 seconds (-1535 ms to +102.4 ms).

Table 1: Dataset Representation Per-class for Fault Type Detection Tasks

Fault Type	# Of Events
100 ms quench	608
3 ms quench	541
electronic quench	674
microphonics	720
heat riser choke	709
control fault	848
single cavity turn off	883

CAVITY FAULT PREDICTION

Cavity fault prediction is a supervised machine learning problem, with ground truth fault labels for recorded data provided by SMEs. The ML model is trained using historical data and used to forecast future events. The dataset used for the prediction task pertaining to normal running signals and faulty events recorded by the data acquisition system. We use the pre-fault signals to predict faults 200 ms before their onset. We propose a two-step pipeline to perform this task which is shown in Fig. 3. Model A is a binary classification network used to distinguish waveforms describing impending faults from stable signals. Those signals identified as an impending fault by model A are used as input to model B which is a multi-class classification network to predict the fault type, all before the fault onset.

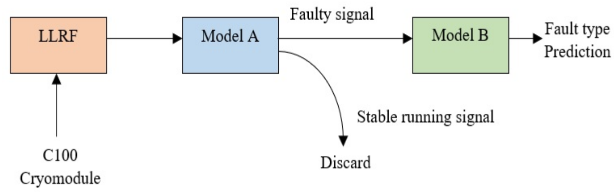

Figure 3: Two-step cavity fault prediction pipeline.

MODEL A: BINARY CLASSIFICATION

The goal of model A is to perform binary classification which identifies waveforms describing impending faults. In the binary classification task, we use U-Net architecture which is a convolutional network consisting of a contracting path (encoder) and an expansive path (decoder) that gives it the characteristic U-shaped architecture [5]. In the encoder, the spatial information is reduced while feature information is increased. The expansive pathway combines the feature and spatial information through a sequence of up-convolutions and concatenations with high-resolution features from the contracting path. The block diagram of the U-Net architecture is shown in Fig. 4. The network is trained to reconstruct the input which consists of normal running examples. During testing, both the normal running signals and impending fault signals pass through the network. Since the network is not trained using the faulty signals, the signal reconstruction loss for the faulty examples will be much larger than the normal running signals. Reconstructions with a higher mean square error (MSE) between input and output are considered a potential fault event.

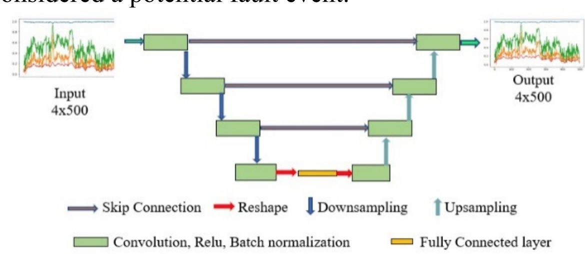

Figure 4: U-Net architecture for binary classification.

We trained the U-Net using 50,000 normal running examples and validate it using 9,000 normal running examples. In the experimental setting we use a learning rate of 10^{-6}, Adam optimizer, MSE loss, and run for 500 epochs. In the test case, we use 997 faulty examples and 997 normal running examples to evaluate the model performance. The receiver operating characteristic (ROC) curve of the model performance is presented in Fig. 5 for different times before the fault onset. As expected, the ROC curve illustrates that model performance improves as the time to the fault onset gets smaller. As the time window of the faulty signals moves away from the fault, the model performance degrades. Table 2 presents the area under the ROC curve (AUC) values for the different times before the fault onset. As the prediction time increases, the AUC value decreases. The AUC value for 200 ms before the fault onset is 0.7126. For 200 ms before the fault onset, we input 1994 test examples to the model from which 886 examples

ISBN: 978-3-95450-232-5 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2022-WEPA29

(with a threshold value of 0.030 for the MSE) identified as impending fault. Among the 886 predicted impending faults, 600 cases were the actual faults.

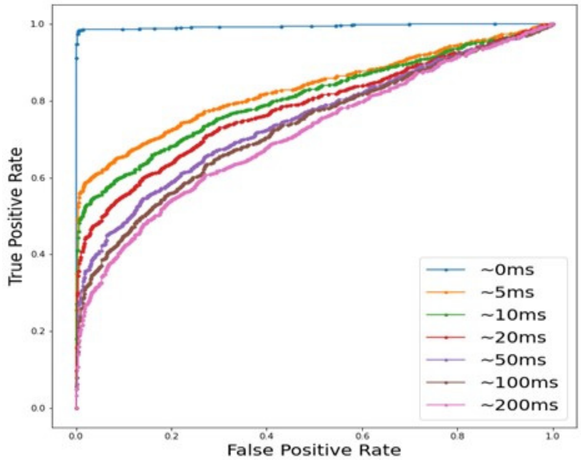

Figure 5: ROC curve for different time before the fault.

Table 2: Area Under the Curve (AUC) of Binary Classification Task at Different Time

Time before fault onset (ms)	AUC
0	0.9936
5	0.8308
10	0.8070
20	0.7803
50	0.7477
100	0.7323
200	0.7126

MODEL B: MULTI-CLASS CLASSIFICATION

The goal of model B is to perform multiclass classification of the different fault types. We use a combination of Long Short-Term Memory (LSTM) and Convolutional Neural Network (CNN) for the model architecture. A schematic of the model is presented in Fig. 6. We concatenate the output of the bi-directional LSTM and 3-layer CNN network to generate the classification output. To reduce overfitting, we use a dropout of 0.8 in the LSTM network. We use 8 classes for this classification task which include 7 different fault types and one class for normal signals. If a normal running signal passes through model A as a faulty signal, model B has the option to identify that normal running signal.

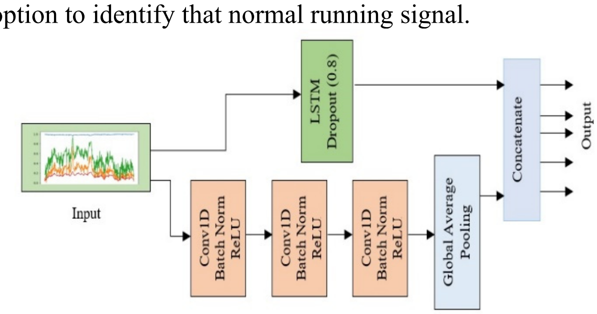

Figure 6: Multi-class classification model diagram.

We train the network using 9,000 examples with a combination of normal and faulty examples. The network validates using 1,994 validation examples. The input of the test dataset for model B is coming from the output of model A which was classified as impending faults. The confusion matrix is presented in Fig. 7. The overall accuracy of the multiclass classification is 76.5% for predictions made 200 ms before the fault onset. Slow-growing faults such as heat riser choke and microphonics showed higher f1-scores (87.3% and 83.3% respectively). Whereas some fast-growing faults such as 3ms quench and single cavity turn off showed lower f1-scores (45.9% and 47.2% respectively). These fast-developing faults only exhibit signatures of an impending event very close to onset.

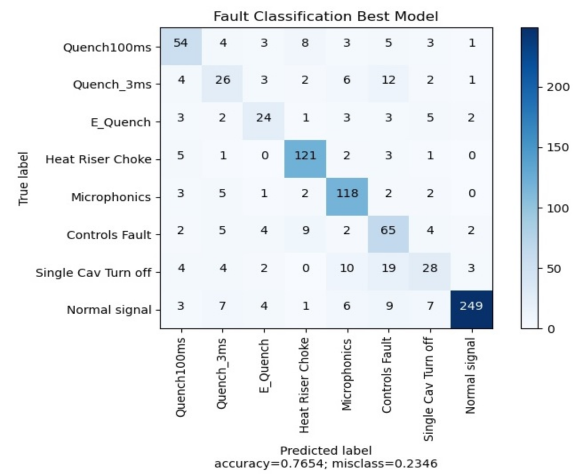

Figure 7: Fault identification confusion matrix.

CONCLUSION

In this work, we have proposed a two-step prediction pipeline for predicting SRF faults in C100 cryomodules. Initial results show the model can predict the fault types 200 ms before the fault onset with reasonable accuracy. The model shows good performance for slow-developing fault types, while identifying fast-developing faults represents a challenge. Future work will explore our ability to make system changes on timescales of a few hundred milliseconds in an effort to mitigate some of the fault that develop over a longer time, such as microphonics.

ACKNOWLEDGEMENTS

The authors would like to thank Tom Powers (Jefferson Lab) for his work in providing labeled data for this work.

REFERENCES

[1] C. E. Reece, “Continuous wave superconducting radio frequency electron linac for nuclear physics research,” *Phys. Rev. Accel. Beams,* vol. 19, no.12, p. 124801, Dec. 2016. doi:10.1103/PhysRevAccelBeams.19.124801

[2] L.S. Vidyaratne, A. Carpenter, R. Suleiman, C. Tennant, D. Turner, Khan Iftekharuddin, and Md. Monibor Rahman, “Initial studies of cavity fault prediction at Jefferson Laboratory”, in *Proc. ICALEPCS'21*, Shanghai, China, Oct. 2021,

ISBN: 978-3-95450-232-5 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2022-WEPA29

paper WEPV025, pp. 700-704.
doi:10.18429/JACoW-ICALEPCS2021-WEPV025

[3] L. Vidyaratne, A. Carpenter, T. Powers, C. Tennant, K.M. Iftekharuddin, M.M. Rahman, and A.S. Shabalina, "Deep Learning Based Superconducting Radio-Frequency Cavity Fault Classification at Jefferson Laboratory," *Front. Artif. Intell*, vol. 4, p. 718950, 2022. doi:10.3389/frai.2021.718950

[4] C. Tennant, A. Carpenter, T. Powers, A. Shabalina Solopova, L. Vidyaratne, and K. Iftekharuddin, "Superconducting radio-frequency cavity fault classification using machine learning at Jefferson Laboratory," *Phys. Rev. Accel. Beams*, vol. 23, p. 114601, 2020. doi:10.1103/PhysRevAccelBeams.23.114601

[5] O. Ronneberger, P. Fischer, and T. Brox (2015), "U-Net: Convolutional Networks for Biomedical Image Segmentation". doi:10.48550/arXiv:1505.04597v1