ISBN: 978-3-95450-247-9 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2024-MOPC44

TRANSFER LEARNING FOR FIELD EMISSION MITIGATION IN CEBAF SRF CAVITIES *

K. Ahammed†, J. Li
Department of Electrical and Computer Engineering
Old Dominion University, Norfolk, VA, USA
A. Carpenter, R. Suleiman, C. Tennant
Thomas Jefferson National Accelerator Facility, Newport News, VA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2024). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

Abstract

The Continuous Electron Beam Accelerator Facility (CEBAF) at Jefferson Lab operates hundreds of superconducting radio frequency (SRF) cavities in its two linear accelerators (linacs). Field emission (FE) is an ongoing operational challenge in higher gradient SRF cavities. FE generates high levels of neutron and gamma radiation leading to damaged accelerator hardware and a radiation hazard environment. During machine development periods, we performed gradient scans to record data capturing the relationship between cavity gradients and radiation levels measured throughout the linacs. However, the field emission environment at CEBAF varies considerably over time as the configuration of the radio frequency (RF) gradients changes and due to the changing behavior of field emitters. An artificial intelligence/machine learning (AI/ML) approach with transfer learning could be a valuable tool to mitigate FE and lower the radiation levels. In this work, we mainly focus on leveraging the RF trip data gathered during CEBAF operations. We develop a transfer learning-based surrogate model for radiation detector readings given RF cavity gradients to track CEBAF's changing configuration and environment. Then, we could use the developed model in an optimization process for redistributing the RF gradients within a linac to minimize radiation levels.

INTRODUCTION

CEBAF (Fig. 1) is a high energy, recirculating continuous wave linear accelerator with two linacs that support four experimental halls [1–3]. These two linacs utilize 400 SRF cavities to accelerate electrons up to 12 GeV through 5.5 passes. Each SRF cavity has a unique gradient threshold above which the phenomenon of field emission occurs [4]. When field emitted electrons are accelerated, they collide with beamline material during their transportation and produce gamma and neutron radiation. This radiation damages CEBAF's equipment, activates components, and contributes to beam trips.

To reduce field emission and lower radiation levels, our earlier effort [5] focused on modelling radiation response for only 32 cavities at the end of the CEBAF's North linac.

* This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under contract DE-AC05-06OR23177
† kaham001@odu.edu

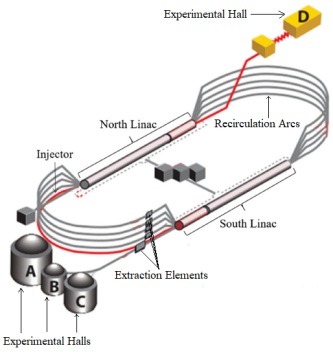


Figure 1: Schematic diagram of CEBAF showing the two linacs. The C100 cryomodules are located at the end of each linac.

That data was gathered was collected during machine development periods. These scans provide snapshot on the behavior of CEBAF, but this behavior changes over time. Existing field emitters could weaken or strengthen and new field emitters could appear. In theory, a model trained on a scan data could be used to reduce FE when paired with an optimization process.

However, that approach becomes less viable as the model's predictive power degrades over time. In this research, we first develop a baseline model for predicting the radiation response of whole North linac. Then we leverage additional scan data and passively collected trip data from normal operations to power a transfer learning approach that helps the model track CEBAF's changing configuration and behavior over time.

MATERIALS AND METHOD

Experimental Set Up and Data Collection

Our initial goal was to focus our efforts on the end of the North linac, where six NDX (neutron dose rate meter with extended capabilities) radiation detectors and four C100-style cryomodules were installed. Typically, a fifth C100 cryomodule resides here, but it had been removed for refurbishment during the time of the study. This area was known for FE activity and had many NDX detectors capable of measuring both neutron and gamma radiation levels. We collected radiation response data for these cryomodules using a combination of single and multi-cavity scan methodologies described in our previous work [5]. We later realized that some cavities from outside this study region were impacting

ISBN: 978-3-95450-247-9 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2024-MOPC44

the radiation readings we collected. This necessitated including cavity gradient data from the rest of the linac, even though we were not actively controlling those cavities. During the study, these passively monitored cavities were subject to maintenance work or other development studies. Thus, they may have unusual characteristics or limited sampling ranges. These data, however, represent the best available data for the entire North linac at the time of the study, and can be supplemented with data passively collected during RF trips. Data were collected this way, separated by roughly eight weeks, in both July and September 2022. These two datasets were produced using similar methodologies, but some elements, such as individual cavity ranges, may vary between them given the changing state of the accelerator. This should not be considered a sampling artifact, but instead a core challenge we must contend with.

In addition to the invasive data collection described above, we also have access to data parasitically collected during RF trips. CEBAF experiences on the order of five beam interruptions due to RF trips per hour in each linac. These interruptions are short in duration (less than five minutes) and involve one or more cavities that experience changing gradients during the trip. Trip data is a valuable source as it tracks changing cavity set points and may capture changing field emitter behavior. However, not all RF cavities will experience trips, and some care is needed when cleaning the individual trips, as the cavities may be poorly regulated during a trip.

To investigate the viability of fine-tuning our model to improve performance over time, we first train a baseline model on our July scan dataset. However, due to the limitations described above we supplement the scan data with July trip data collected for two weeks after the scan data to train the baseline model. We then experiment with fine-tuning the baseline model in two ways. First, we fine-tune the baseline model on trip data collected over the two-week period surrounding the September scan dataset. Secondly, we fine-tune the baseline model on a large portion of the September scan dataset. Both fine-tuned models and the baseline model are evaluated on a withheld portion of the September scan dataset. The differences in performance of these models on the withheld September data give good indications of the viability of fine-tuning, despite the stated challenges in our datasets.

Dataset and Data Preprocessing

The data used in this research are from both gradient scan and RF trip datasets (Fig. 2). The gradient scans cover 32 actively studied cavities (cryomodules 1L22 -1L25) and 160 passively studied cavities (cryomodules 1L02 – 1L21). RF trip data includes cavities that transition through a range of gradients while recovering. Typically, RF trips only involve a small number of cavities out of the entire linac and many cavities will not trip over a few weeks (Fig. 3). The data consists of two classes, cavity gradients (MV/m, input) and radiations (rem/h, output), where radiations are further categorized into neutron and gamma radiations. We normalize

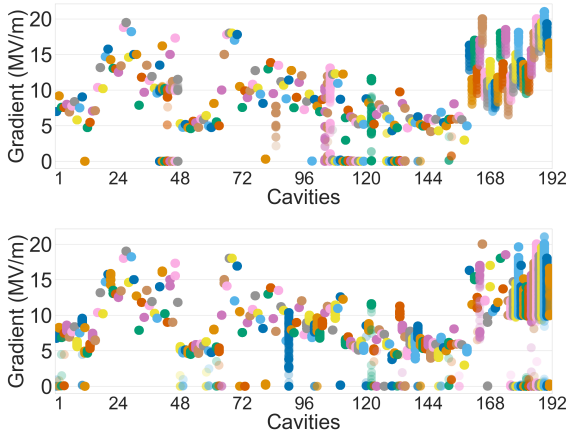


Figure 2: Cavity gradient distribution of raw July scan dataset (top) and raw September trip dataset (bottom). Each point represents a single observed gradient for that cavity.

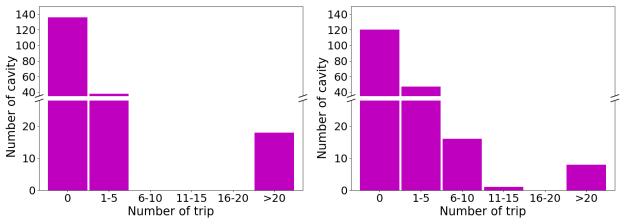


Figure 3: Number of cavities with trips for July trip dataset (left) and September trip dataset (right).

both input and output of training dataset using Z-score normalization. To normalize the test dataset, we use the mean and standard deviation of the training dataset.

The gradient scan datasets are both heavily curated through their collection process and require only minor refinement. The trip data set requires more extensive cleaning to limit the range of gradient changes to relevant values and exclude poorly regulated cavity behavior. This cleaning is most effectively achieved by only keeping examples of cavity gradients where all cavities are within 3 MV/m of their setpoint.

As previously stated, the dataset used in training a baseline model is a combination of July scan data and July trip data. After cleaning, the July trip data set was considerably larger than the July scan data set. In order to keep the data balanced we randomly sample 10% of each trip for use in training. This results in 1,278 trip datapoints from 573 trips that are combined with 1,244 scan datapoints from 311 unique set point combinations. This baseline dataset is randomly split by group shuffling into train/validation/test set using 60%/20%/20% splits.

The first fine-tuning dataset is comprised of trip data collected from one week prior to the September scan set and one week after. After pre-processing, this dataset includes 3,431 datapoints from 65 trips. The second fine-tuning dataset is made up of 75% of the September scan dataset and includes 1,107 datapoints. The remaining 25% of the September scan set is withheld for testing and contains 369 datapoints of 92 unique set point combinations.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2024). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-247-9 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2024-MOPC44

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2024). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

Model Architectures

We use a convolutional neural network (Fig. 4) consisting of one-dimensional kernels for our baseline model. We train this model on the baseline dataset from July and evaluate the performance of the pre-trained model on the September test set. We then conduct two experiments to evaluate the performance improvement through fine-tuning. In experiment 1, we use the September trip dataset to fine-tune the model. In experiment 2, we fine-tune the model on September scan data. During fine-tuning, we train the pre-trained model without freezing any layers. Finally, we evaluate the fine-tuned models on the withheld September test set.

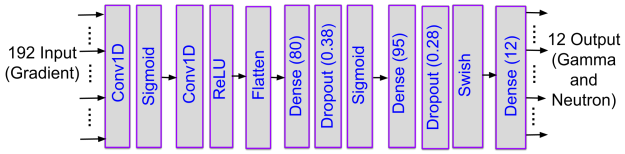


Figure 4: Baseline model architecture of the convolutional neural network (CNN).

RESULTS AND DISCUSSION

As explained above, we first develop a baseline CNN model for transfer learning approach for modelling radiation in the whole North linac. This baseline model performs well on the July test data, but exhibits a marked reduction on the September test data (Table 1). This is consistent with our observations that CEBAF's RF configuration and field emission response change over time. However, we find mixed results from fine-tuning. Fine-tuning on the September scan dataset improves the performance of the model considerably, while fine-tuning on the September trip data helps in some cases and hurts in others (Table 1 and Fig. 5).

In general, the September scan set should provide the upper bound for the potential performance improvement of fine-tuning. This assumption appears to hold as the model fine-tuned on it provides the best performance in general, despite weaker performance on the radiation detected at 1L22.

Table 1: Model performances on both a withheld test set (20% of combination of July scan and July trip set) similar to their training data and on the withheld September test set (25% of September scan set).

Models	Metrics	July Test	Fine-tuning Validation	September Test
BL CNN	MSE	4.474		20.014
	MAE	1.036		2.439
	MAPE	8.5%		19.8%
Trip FT CNN	MSE		5.940	14.143
	MAE		1.086	1.928
	MAPE		31.1%	14.5%
Scan FT CNN	MSE		2.162	1.711
	MAE		0.769	0.700
	MAPE		6.5%	6.6%

BL CNN = Baseline convolutional neural network, FT = Fine-tuning, MSE = Mean squared error, MAE = Mean absolute error, MAPE = Mean absolute percentage error.

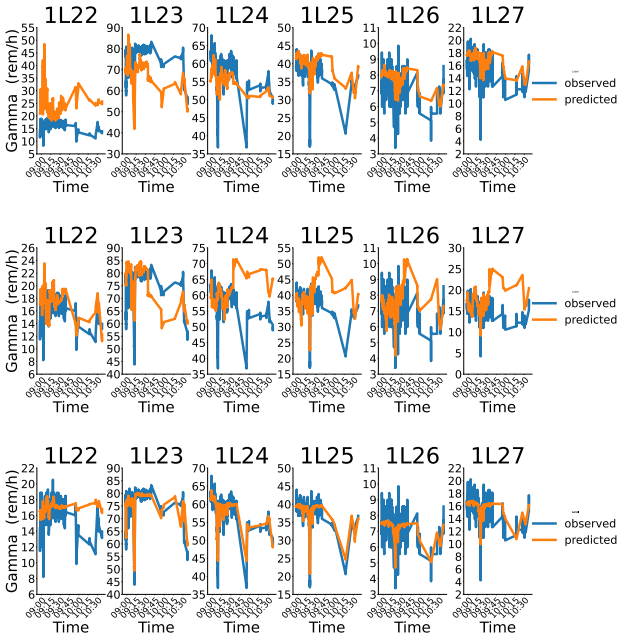


Figure 5: Observed and predicted time series plots for gamma radiation by the baseline CNN model (top), trip fine-tuned CNN model (middle), and scan fine-tuned CNN model (bottom) on September test dataset.

The performance of the fine-tuned model on the scan set shows the feasibility of fine-tuning given a semi-representative dataset. However, the model fine-tuned on trip set shows less improvement overall. One major challenge with using the trip data is that some cavities may experience few, if any, trips. This provides little information on the impact of that cavity. For example, not all C100s experienced trips and only a few of the non-C100 cavities had a significant number of trips (~ 19) that likely contain information on upstream detector responses. One possible mitigation is to supplement the trip data with scan data to better represent all cavities during model training.

While the trip FT CNN model performs worse on several detectors, it appears to perform the best for 1L22. This is likely because non-C100 cavities were better sampled in trip data than in the September scan data. This shows that given the right trip data, fine-tuning on the parasitic trip data can meaningfully improve model performance similarly to training on new invasively acquired scan data.

CONCLUSION

In this research, we developed a baseline CNN model and demonstrated the ability to use a transfer learning approach on data collected both invasively and non-invasively for improving radiation predictions over time.

We are currently working to improve our data collection through studies on synthetic data and real-world sampling methods that cover the entire linac. We are also pursuing the use of offline optimization to generate gradient distributions that reduce radiation in a linac.

ISBN: 978-3-95450-247-9 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2024-MOPC44

REFERENCES

[1] M. F. Spata, "12 GeV CEBAF Initial Operational Experience and Challenges", *in Proc. IPAC'18*, Vancouver, Canada, Apr.-May 2018, pp. 1771-1775. doi:10.18429/JACoW-IPAC2018-WEYGBD1

[2] A. Carpenter *et al.*, "Using AI for Management of Field Emission in SRF Linacs", *in Proc. ICALEPCS'21*, Shanghai, China, Oct. 2021, pp. 970-974. doi:10.18429/JACoW-ICALEPCS2021-THPV043

[3] 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

[4] R. L. Geng, A. Freyberger, R. A. Legg, R. Suleiman, and A. S. Fisher, "Field Emission in SRF Accelerators: Instrumented Measurements for Its Understanding and Mitigation", *in Proc. IBIC'17*, Grand Rapids, MI, USA, Aug. 2017, pp. 470-477. doi:10.18429/JACoW-IBIC2017-TH1AB1

[5] K. Ahammed *et al.*, "Field Emission Mitigation in CEBAF SRF Cavities Using Deep Learning", *in Proc. NAPAC'22*, Albuquerque, NM, USA, Aug. 2022, pp. 676-678. doi:10.18429/JACoW-NAPAC2022-WEPA25

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2024). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.