

MACHINE LEARNING BASED PREVENTIVE MAINTENANCE AND AUTONOMOUS POWER CONTROL FOR RF CAVITIES IN A FREE-ELECTRON LASER

T. Yuvaraj*, Columbia University, New York, United States
A. Sharma, Indian Institute of Technology Delhi, New Delhi, India
B. K. Sahu, Inter-University Accelerator Centre, New Delhi, India

Abstract

This paper presents a machine learning system for automated RF power control during high-power conditioning of the RF photocathode gun at the Free Electron laser based Delhi Light Source (DLS) facility, IUAC. A hybrid model combining K-means clustering, linear regression, and a decision tree classifier categorizes ion-pump current ranges, predicts baseline behavior, and detects anomalies in real time. Integrated with the EPICS control system, the algorithm adjusts power every 50 ms to prevent trips and recover from faults, enabling effective automated conditioning with reduced operator workload.

INTRODUCTION

The Delhi Light Source (DLS) at IUAC, New Delhi, is a compact, pre-bunched Free Electron Laser facility designed to produce electron beams up to 8 MeV using a 2.6-cell normal-conducting RF photocathode gun at 2860 MHz [1,2]. The RF gun is powered by a high-power klystron modulator based RF system from M/S Scandionova Systems [3] (Fig. 1), delivering peak RF power of several megawatts [4,6]. RF conditioning is required to reach the target accelerating gradient while minimizing breakdown and dark current and keeping stable ultra-high vacuum conditions ($< 10^{-9}$ mbar) for reliable operation, particularly when using semiconductor photocathodes.


Before beam production, the gun must undergo RF conditioning, gradually increasing power while maintaining stable vacuum. Breakdowns and multipacting cause pressure rises that trip the system. This manual process requires continuous monitoring, with nighttime operation at IUAC making sustained oversight impractical. A predictive ML system to automate RF conditioning, reducing operator workload and trip frequency, was developed. In order to protect the source against high reflected power during conditioning, an SF₆-based isolator was installed [5], enabling conditioning to higher field gradients up to 100 MV/m.

Conditioning Challenges


High-power RF conditioning of metallic cavities is inherently challenging due to the risk of RF breakdowns, multipacting, and vacuum degradation. The process requires:

- Starting from low power and gradually increasing, with careful monitoring of vacuum levels.

- Continuous tracking of the gun resonance frequency with the signal generator to minimize reflected power.
- Optimization of the photocathode plug position for impedance matching.
- Maintaining VSWR interlock conditions within the klystron's specified limits.



(a) RF source side: ScandiNova modulator, klystron, and IPC01 ion pump.



(b) Load side: RF photocathode gun and IPC07 ion pump.

Figure 1: RF system for the photocathode gun at IUAC-DLS, showing (a) the source side with the klystron modulator and (b) the load side with the photocathode gun and the IPC07 ion pump whose current is the primary input to the ML model.

DATA COLLECTION AND EXPLORATION

A Python client interfacing with the EPICS (Experimental Physics and Industrial Control System) framework logs process variables to CSV files every 2 s, with accelerated sampling during periods of high-current fluctuation. The key process variables monitored include: machine operating state (integer-encoded as 13=Trigger, 9=High Voltage, 5=Standby, 1=Off), high-voltage reading, pulse width, ion-pump pressures and currents at three locations (IPC01,

* tushnim.yuvaraj@columbia.edu

klystron; IPC03, circulator; IPC07, cavity), forward and reflected RF output power, driver amplifier forward power, and the commanded set power and frequency.

Among these, the IPC07 current (directly proportional to cavity pressure) and the set power are the most critical variables. The entire control algorithm operates by adjusting the set power in response to IPC07 current fluctuations, while the RF Forward power (RFOutFwd) dropping to zero serves as the definitive indicator of a system trip, which is what the algorithm aims to prevent.

Several analyses guided model development. A correlation matrix computed across all process variables using Pandas showed limited inter-variable correlation for anomalous events, partly because the 2 s sampling interval is too slow to capture the fast rising edge of current surges during RF breakdowns.

Data exploration during manual conditioning revealed that the t-distributed Stochastic Neighbor Embedding (t-SNE) algorithm, applied for dimensionality reduction, produced clear clustering by set power level (Fig. 2), confirming RF input power as the dominant organizing parameter of the system state. Each data point in the t-SNE map represents a single timestep snapshot of all monitored process variables. The algorithm, which has no access to the set power variable during embedding, maps these high-dimensional points to two dimensions while preserving local neighborhood structure. The emergence of distinct, well-separated clusters corresponding to each power setting indicates that the full system state (vacuum levels, ion-pump currents, reflected power, and other variables) reorganizes coherently at each operating power. No comparable clustering was observed when coloring by IPC07 current or any other single variable, indicating that anomalous IPC07 excursions are not a distinct system state but rather deviations within a power-level regime. This result directly motivates a power-conditional approach to anomaly detection: a single global threshold for IPC07 current would produce false alarms at high power (where the baseline is naturally elevated) or miss genuine precursors at low power.

Analysis of the mean baseline and standard deviation of IPC07 current at each power level showed a clear upward trend with increasing power, quantitatively confirming the power-dependent behavior revealed by t-SNE and providing the foundation of the predictive model.

THE MODEL

The deployed model has three sequential stages: unsupervised clustering for range identification, supervised regression for baseline prediction, and supervised classification for real-time anomaly severity assessment.

K-Means Clustering was applied to the IPC07 current data to identify natural groupings. The algorithm partitions observations into K clusters by minimizing the within-cluster




Figure 2: t-SNE visualization of process variable space colored by set power level (in dBm). The horizontal and vertical axes (tsne-2d-one, tsne-2d-two) are the two embedding dimensions produced by the t-SNE algorithm and carry no physical units. Each point represents a single timestep snapshot of all monitored variables. Distinct clusters form by operating power.




Figure 3: K-means clustering of IPC07 current data showing the four identified operational ranges. The log-scale vertical axis reflects the highly skewed distribution, with the majority of observations in the Safe range.

sum of squares:

$$\arg \min_S \sum_{i=1}^K \sum_{\mathbf{x} \in S_i} \|\mathbf{x} - \boldsymbol{\mu}_i\|^2 \quad (1)$$

where S_i is the set of observations in cluster i and $\boldsymbol{\mu}_i$ is the cluster centroid. The analysis identified four operational ranges (Fig. 3): **Safe** (0), **Moderate** (1), **High** (2), and **Very High** (3), corresponding to increasing deviation from the baseline current. The boundaries between ranges were determined by examining the maximum and minimum current values within each cluster.

Linear Regression predicts the expected baseline IPC07 current and its standard deviation at each power level. These predictions enable classification of any observed current value into the four K-means-derived ranges relative to the current operating power, capturing the observed upward trend in both baseline current and its variability with increasing RF power.

Decision Tree Classifier takes the current range classification along with other process variables as input and outputs a severity level (0–3) indicating the likelihood of an impending trip. Decision trees were chosen for their interpretability, which allows operators to understand the classification logic, and their fast inference time, which is

critical for the 50 ms control loop. They also handle non-linear decision boundaries without requiring feature scaling.

CONTROL PROGRAM AND DEPLOYMENT

The trained model is integrated into a Python control program interfacing with EPICS. The user specifies the target conditioning power range at startup. Every 50 ms, the program reads process variables, classifies the IPC07 current using the regressor and K-means-derived ranges, obtains a severity level from the decision tree, and executes the corresponding control action (Table 1).

Table 1: Control Actions Based on Severity Level

Level	Action
0 (Safe)	No change; after 5 min stability, increase power by set increment
1 (Moderate)	Reduce power by small decrement
2 (High)	Reduce power by larger decrement
3 (Very High)	Shut off RF; restart at reduced level

The program includes automatic trip recovery: if RFOut-Fwd drops to zero, it waits for vacuum to recover and assesses the system health. If a serious interlock has occurred, it stops, waiting for operator intervention; otherwise, it restarts conditioning at a reduced power level. All power adjustments and system events are logged with timestamps for post-analysis. During stable operation, the algorithm steadily increases power in set increments after sustained safe conditions; when anomalies are detected, it autonomously reduces the set power to prevent trips.

The software stack uses `pyepics` for EPICS communication, `scikit-learn` for the K-means, linear regression, and decision tree models (serialized via `joblib` for deployment), and `Pandas/NumPy` for data handling. Model retraining is supported through a dedicated script that accepts new CSV data files and regenerates all model components.

RESULTS AND FUTURE WORK

The system has been deployed at the IUAC FEL facility and demonstrates effective trip prevention during unsupervised conditioning runs. It autonomously adjusts power in response to vacuum excursions and recovers from faults without operator intervention. Nighttime operator presence requirements have been significantly reduced, allowing conditioning to proceed for extended periods.

With the present RF system, the RF photocathode gun has been conditioned to a field gradient of 100 MV/m for the desired 4 μ s pulse length, with electron beam energy of more than 6.5 MeV verified at 10 MW forward power [5]. Conditioning up to desired power level of more than 12 MW has been achieved for the 8 MeV beam energy target, which would have taken several manhours to achieve otherwise.

The limitations of occasionally overly conservative power reductions for minor fluctuations, and limited capture of fast transients at 2 s sampling are being worked out to improve by

reinforcement learning for adaptive policy optimization [7], higher sampling rates enabling LSTM-based models [8], and transition to online learning.

CONCLUSION

A predictive ML system combining K-means clustering, linear regression, and decision tree classification has been deployed for automated RF conditioning of the IUAC-DLS photocathode gun. The system operates at 50 ms intervals within the EPICS control framework, classifying IPC07 current deviations relative to power-dependent baselines and executing graduated control actions to prevent system trips. Deployed results demonstrate effective unsupervised conditioning with significantly reduced operator workload, achieving field gradients of 100 MV/m and beam energies exceeding 6.5 MeV. Future development will focus on reinforcement learning for adaptive policy optimization and higher-rate data acquisition to enable sequence-based models.

ACKNOWLEDGEMENTS

The suggestions and technical inputs received from members of the FEL group of IUAC Dr. S. Ghosh, Mr. P. Patra, Mr. B. Karmakar, Mr. J. Karmakar and Ms. Madhuri for the successful implementation of the automated RF conditioning project, are highly acknowledged. The authors acknowledge the contribution of accelerator supporting groups of IUAC and contractual persons responsible for RF conditioning for their support during this project work. The encouragement and support received from Prof. A.C. Pandey, Director, IUAC is acknowledged.

REFERENCES

- [1] S. Ghosh *et al.*, “Status of the development of Delhi Light Source (DLS) at IUAC,” *Nucl. Instr. Meth. B*, vol. 402, pp. 358–363, 2017. doi:10.1016/j.nimb.2017.03.108
- [2] A. Deshpande *et al.*, “Experimental results of an rf gun and the generation of a multibunch beam,” *Phys. Rev. ST Accel. Beams*, vol. 14, p. 063501, 2011. doi:10.1103/PhysRevSTAB.14.063501
- [3] Scadinova Systems, <https://scadinovasystems.com>
- [4] B. K. Sahu *et al.*, “High Power RF Conditioning of the Photocathode Gun for IUAC_DLS,” in *Proc. Indian Particle Accelerator Conf. (InPAC’22)*, VECC, Kolkata, India, Mar. 2022.
- [5] B. K. Sahu *et al.*, “Final Commissioning of the High-Power RF System for Conditioning of the RF Photocathode Gun at Higher Field Gradient,” in *Proc. Indian Particle Accelerator Conf. (InPAC’23)*, 2023.
- [6] B. K. Sahu *et al.*, “RF System for the Free Electron Laser Based Delhi Light Source,” in *Proc. Indian Particle Accelerator Conf. (InPAC’19)*, IUAC, New Delhi, India, Nov. 2019.
- [7] X. Pang, S. Thulasidasan, and L. Rybarczyk, “Autonomous Control of a Particle Accelerator Using Deep Reinforcement Learning,” 2020. doi:10.48550/arXiv.2010.08141
- [8] D. Gilardi *et al.*, “Application of LSTM Neural Networks for RF Cavity Conditioning,” in *Proc. IPAC’21*, MOPAB344, 2021. doi:10.18429/JACoW-IPAC2021-MOPAB344